MASARYK UNIVERSITY
FACULTY OF INFORMATICS



Large-Scale Molecular Dynamics Simulations for Highly Parallel Infrastructures

PH.D. THESIS PROPOSAL

Jana Pazúriková

Brno, 2014

arXiv:1402.7216v1 [cs.DC] 28 Feb 2014

Adviser: prof. RNDr. Luděk Matyska, CSc.

Adviser's signature: _____

Hereby I declare that this thesis proposal is my original authorial work, which I have worked out by my own. All sources, references, and literature used or excerpted during preparation of this work are properly cited and listed in complete reference to the due source.

I would like to thank prof. Matyska and dr. Křenek for their supervision and consultation. The challenging task that led to this dissertation proposal has been suggested by dr. Vácha and dr. Kulhánek, I also greatly appreciate their advice on chemical aspect of work. Last but not least, I thank my colleagues from Sitola laboratory for their valuable comments.

Contents

1 Introduction 1
2 State of the Art 4
2.1 Molecular Dynamics Simulations in a Nutshell 4
2.2 Calculation of Long-Range Interactions 9
2.3 Multilevel Summation Method 10
2.4 Acceleration of Molecular Dynamics Simulations . . 14
2.5 Parallel-in-Time Computation 18
2.6 Parareal Method 20
2.7 Limitations of the Parareal Method 22
2.8 Summary 22
3 Research Questions and Proposed Solutions 24
3.1 Research Questions 24
3.2 Proposed Solutions 25
3.3 Complexity and ability to run in parallel 27
3.4 Convergence 29
3.5 Future work 30
3.6 Evaluation 30
4 Aim of the Work 32
4.1 Expected Results 32
4.2 Schedule 34

Glossary **35**

Acronyms **37**

1 Introduction

Computational chemistry allows researchers to experiment *in sillico*: by running a computer simulations of a biological or chemical processes of interest. Computer models of the chemical processes offer higher resolution at the cost of decreased accuracy. Molecular dynamics (MD), a tool of computational chemistry, evaluates movements of particles caused by their interactions, it examines changes in time. Molecular mechanics (MM) models treat atoms as points with mass and charge and approximate their interactions with empirical functions. Molecular dynamics with molecular mechanics (MM MD) simulates N-body problem of atoms: it computes movements of atoms according to Newtonian physics and empirical descriptions of atomic interactions. Within each step the computer evaluates the forces exerted upon each atom caused by bond stretching, angle bending, torsion bending, van der Waals and electrostatic potentials and then moves the atoms accordingly. In order to capture the fastest oscillations occurring at atomic scale, the vibrations of bonds containing the hydrogen, the timestep of integration scheme is $\sim$ 1 femtosecond (10^{-15} s). Common experiments simulate tens of thousands of atoms for hundreds of nanoseconds. [Jen07, Lew10]

However, many interesting processes occur at longer timescales: tens of microseconds and more. Their simulations include tens, hundreds, thousands of billions of steps, each of them computationally demanding due to the evaluation of electrostatic interactions. Therefore, MM MD simulations need high performance computing resources and approaches. Long wallclock time impedes the research in areas that verify various proposed solutions to their research problems through MD simulations before using experiments, such as drug discovery or nanomaterial development.

Several different approaches deal with computational demands of MM MD. First, methods such as coarse-grained modelling [RB98], discrete MD [PDD11] and all algorithms for evaluation of electrostatic interactions [Koe06] further simplify the model and approximate its parts. Second, low level acceleration through specialized hardware [SDD+07] and GPU [SPF+07, ALT08, LSVMW08, HSS09] achieves high speed-up compared to computation on CPU, however, at increased cost. And last, computation on parallel and distributed infrastructures has been able to cut the time to result thanks to spatial domain decomposition[HKvdSL08, CCD+05, PBW+05, CHB+05, Pli03, ABF+13]. Unfortunately, with growing amount of computational resources, issues with scaling have arisen.

Current parallel implementations of MM MD algorithms exhibit almost perfect weak scaling, i.e. they are able to simulate even large systems at reasonable wallclock time, if they are provided with enough computational power. However, the fixed-size simulation will not run faster if provided with arbitrary number of computational resources, the strong scalability hits

the wall. After reaching some critical limit in number of computing cores, adding more will not shorten the time to result. The fine granularity of the problem per core will result in high communication and synchronization overhead. The state-of-the-art MM MD algorithms with the highest strong scaling can saturate up to half a million cores [RGC+09, AYF+13]. With the dawn of exascale computers, we want to shift the level of parallelism further.

As the spatial decomposition does not suffice, we propose rather uncommon approach: to calculate MM MD parallel-in-time. Current research in molecular dynamics simulations is beginning to focus on time parallelism. Long simulations of protein folding have been conducted at highly distributed infrastructure through Copernicus and Folding@Home projects [PLP+11, LSSP02]. These frameworks exploit coarse-grained time parallelism by many parallel short simulations that explore the conformation space of the protein. Yu et al. [YSC06] use data from previous similar simulations to predict system states at future time points. Parallel-in-time methods, such as [Nie64, ML67, VvdV94, LMT01], calculate the results of a time dependent differential equation in several successive time points simultaneously. This fine-grained form of time parallelism without *a priori* knowledge is almost unknown in MM MD. Baffico et al. [BBM+02] wrote the first, rather limited, publication on this topic in 2002. They concluded that parallel-in-time calculation "can be very useful". In 2013, Bulin wrote a master thesis [Bul13] where two parallel-in-time algorithms are compared in molecular dynamics simulations.

One of the parallel-in-time methods, the parareal method, first approximates the results ahead a few timesteps with less accurate but cheap coarse function $\mathcal{G}$ and then iteratively corrects the results in parallel with accurate but expensive fine function $\mathcal{F}$. The appropriate choice of coarse function determines the convergence and speed-up of the method. We have analyzed several combinations of fine and coarse methods and found two that promise high theoretical speed-up and reasonable convergence. We have designed the parareal multilevel summation method with simple cutoff method or Wolf summation method as the coarse functions and multilevel summation method as the fine function.

The aim of the dissertation is to study and develop algorithms that would enable us to simulate large systems for long simulation times. We want to achieve it by incorporating the time parallelism into the calculation of long-range interactions in MM MD simulations. We will study and implement the parareal multilevel summation method; evaluate its correspondence with results from experimental chemistry and compare the accuracy, speed-up and scalability with the best implementations of other methods.

This thesis proposal continues with three more sections. First we will overview the area of molecular dynamics, describe the issues with the calculation of long-range interactions, review several improvements that ac-

celerate MD and introduce parallel-in-time computation. We will present state-of-the-art methods that are able to utilize large number of processors to speed up long simulations. In the second chapter, we will state the research questions and propose solutions. We will introduce the novel parareal multi-level summation method, analyze the theoretical speed-up and convergence and suggest several directions of future work. In the third chapter, we will state the aim of the work and schedule future progress.

2 State of the Art

2.1 Molecular Dynamics Simulations in a Nutshell

Introduction

Molecular dynamics evaluates the movements of particles caused by their interactions. Many models describe the interactions: molecular mechanics, quantum mechanics, coarse-grained models and more. This thesis proposal deals with simulations of molecular dynamics with molecular mechanics model and the term *molecular dynamics* (MD) will be used in this sense from now on if not specified otherwise. Through MD simulations, researchers can observe and experiment with a model of molecules in order to understand chemical or biological processes and predict macroscopic properties by detailed knowledge of atomic movements caused by their interactions [Jen07]. Bonded atoms interact due to bond stretching, angle bending and dihedral torsions. Non-bonded atoms interact due to van der Waals and Coulomb interactions. Coulomb (also called electrostatic) interactions range to long distances, therefore they should be calculated between every pair of atoms. The evaluation of long-range interactions with current methods has $\mathcal{O}(NlogN)$ or $\mathcal{O}(N)$ asymptotic complexity, where N is the number of atoms. Still it remains the most demanding part of the calculation. Interactions can be described by the potential function between the particles and the force caused by this potential changes the positions of particles. The movement is governed by Newton's second law of motion commonly known as $\mathbf{F} = m\mathbf{a}$ where $\mathbf{F}$ is the force, m is the mass of the atom, $\mathbf{a}$ is the acceleration of the atom. Each step of the simulation (that takes usually 2 femtoseconds of the simulation time), the potential of interactions is calculated, Newton's second law of motion in form of the partial differential equation is numerically solved and positions of atoms are updated. First computer simulations of molecular dynamics were published in 1950s [AW59]; since then, many methods have been developed to accelerate the calculation. Further approximations of electrostatic interactions have been designed and the simulations are calculated on parallel and/or distributed infrastructure.

Motivation

Thanks to molecular dynamics simulations, chemists and biologists can better understand how processes happen on atomic scale. That helps them for example:

1. to simulate how a protein folds or examine conditions under which it misfolds that can shed some light to the cause of Alzheimer's disease or cancer [KM02, SNPG02, LH11];

2. to see how enzyme interacts with a substrate and how a drug interferes with this interaction [BdOD+13, GaOdC+11, SKS11];

3. to see how designed nanomaterial interacts with other substances or reacts to the external force that can foretell its properties [LL12, BBSD13, TP10];

4. to understand how phenomena occur and explain underlying physical reasons [OCP+13, ZPY+13, SS07].

MD simulations differ in the size (number of atoms), simulation time (number of steps $\times$ length of single timestep). These characteristics and the aim of simulations determine what further approximations or improvements can be done without major negative influence.

Researchers may want to explore the trajectory of each atom step-by-step. In that case, classical MD simulation provides the data. However, various force fields (the set of parameters that determines the calculation of atomic interactions) have been developed for various types of molecules—inorganic/organic, specialized/general [CCB+95, MBF01, Hal96, RCC+92]. Also, long simulations of small systems have different issues than short simulations of large systems when implemented for parallel or distributed infrastructure.

Researchers may want to identify different conformations (3D structure of a molecule in energy minimum) that the molecule can get to from the initial state and explore the conformation space. Or they want to assess the energy surface (the function that maps the coordinates of atoms to the energy of the system). In these cases, they can accelerate the computation by adding the artificial, biased potential [LG08, TV77, MF04, SO99]. The simulated system would not behave as it does in nature, however, unbiased properties can be reconstructed.

Molecular dynamics as a tool of computational chemistry helps scientists in many fields and the wallclock time of the simulation crucially influences the phenomena they can study.

Simulation loop

A general simulation algorithm first takes input data—types of atoms, the topology, partial charges q_i, positions $\mathbf{r}_i$ and velocities $\mathbf{v}_i$ and several other parameters for all particles in the system—and then iteratively repeats the following steps:

1. calculate the potential and forces

$$\begin{gathered} U_{all} = U_{bonded} + U_{non-bonded} = \\ U_{bond} + U_{angle} + U_{torsion} + U_{Waals} + U_{Coulomb} \end{gathered} \tag{1}$$

where U_* is the potential due to $*$;

$$\mathbf{F}_i = -\frac{\partial U_{all}}{\partial \mathbf{r}_i} \tag{2}$$

where $\mathbf{F}_i$ is the force exerted on atom i;

2. move particles

$$\mathbf{F}_i = -\frac{\partial U_{all}}{\partial \mathbf{r}_i} = m_i \mathbf{a}_i = m_i \frac{\mathbf{v}_i'}{dt} = m_i \frac{\mathbf{r}_i''}{dt^2}; \tag{3}$$

3. update time, optionally generate output. [Jen07]

The output of the simulation includes the trajectories of particles, forces and energy of the system. Properties of the system such as an average potential energy or the viscosity of the liquid are processed from output data by using further physical and chemical equations and applying statistical methods.

The differential equation (3) is solved by numerical methods, most common are leap-frog [AT89] and velocity Verlet [SABK82] integration schemes. Leap-frog scheme, a modification of Verlet method, starts with positions $\mathbf{r}$ at time t and velocities $\mathbf{v}$ at time $t + \frac{1}{2}\Delta t$ where Δt is the timestep. Then it updates the positions according to velocities, evaluates the potential and acquires acceleration in the next time point and updates the velocities [Lea01]:

$$\begin{aligned} \mathbf{r}(t + \Delta t) &= \mathbf{r}(t) + \Delta t \mathbf{v}(t + \tfrac{1}{2}\Delta t) \\ \mathbf{r}(t + \Delta t) &\rightarrow U(\mathbf{r}(t + \Delta t)) \rightarrow \mathbf{a} \\ \mathbf{v}(t + \tfrac{3}{2}\Delta t) &= \mathbf{v}(t + \tfrac{1}{2}\Delta t) + \mathbf{a}(t + \Delta t)\Delta t \end{aligned} \tag{4}$$

Velocity Verlet starts with position and velocities at time t and continues [Lea01]:

$$\begin{aligned} \mathbf{v}\left(t + \tfrac{1}{2}\Delta t\right) &= \mathbf{v}(t) + \tfrac{1}{2}\mathbf{a}(t)\,\Delta t \\ \mathbf{r}(t + \Delta t) &= \mathbf{r}(t) + \mathbf{v}\left(t + \tfrac{1}{2}\,\Delta t\right)\,\Delta t \\ \mathbf{r}(t + \Delta t) &\rightarrow U(\mathbf{r}(t + \Delta t)) \rightarrow \mathbf{a} \\ \mathbf{v}(t + \Delta t) &= \mathbf{v}\left(t + \tfrac{1}{2}\,\Delta t\right) + \tfrac{1}{2}\,\mathbf{a}(t + \Delta t)\Delta t \end{aligned} \tag{5}$$

More accurate methods, such as Runge-Kutta, would require the computationally demanding evaluation of the forces a few times within one step and therefore they can not be applied.

The values for input data come from various sources. The structure of many proteins, viruses and other chemical compounds can be found in repositories such as [BHN03]. Partial charges of atoms are determined by the force field. The initial velocities are randomly assigned by Maxwell-Boltzmann distribution corresponding to a low temperature. In a short simulation, the system is heated up to usually 300 K by adding random

kinetic energy to atoms. Often, similar short simulations follow to introduce the solvent, minimize the energy and stabilize the pressure. After them, the system is prepared for the main simulation, relaxed and with stable pressure and temperature.

Issues and Limitations of Molecular Dynamics

Molecular dynamics has several issues, some caused by the characteristics of simulated chemical processes, most of them caused by the nature of the method itself.

In models of computational chemistry, the energy surface maps coordinates of atoms to the energy. This function has many local minima that correspond to more or less stable states of the system. In the interesting processes the system usually crosses the energy barrier and transforms from one state to another. However, crossing the barrier occurs with probability that exponentially relates to its height, i.e. the higher barrier, the less probable crossing. Therefore, classical MD simulations sometimes have to simulate for long simulation time for crossing of the energy barrier to happen. As it can be rather difficult and lengthy to cover and sample whole energy surface, the issue is called the sampling problem [Jen07].

Needed long simulation times directly lead to long wallclock times as MD solves the initial value problem in a sequence of steps and, moreover, the integration scheme has small timestep due to high oscillations of bonds that contain the hydrogen. The evaluation of the potential between atoms in each step remains computationally demanding despite many approximations of long-range interactions. Moreover, evaluation of long-range interactions requires communication between all processors calculating spatially decomposed parts. Many decomposition techniques rather compute the same values on two different processors than send a message which stresses the high temporal cost of communication [BDS05].

The relative error of integration scheme that affects the forces achieves 1% for common 2 fs timestep[Mat13]. This would cause unwanted changes in pressure and temperature. Therefore, MD simulation softwares regularly (every few steps) check the temperature and pressure and alter the atom velocities or the volume to stabilize them. The simulations keep the number of atoms, the temperature and the pressure or the volume constant.

Despite errors due to rounding off, integration scheme, force field and approximations of long-range interactions, MD is considered to be accurate [ST05]. The probability density of the system states and global quantities of the simulated system correspond with results acquired by experiments. Sometimes, the simulated system blows up [Gro13]. Blowing up refers to the state with extremely large force that causes the failure of the integrator. The reasons the system gets into such state include insufficient prior energy minimization, large timestep, inappropriate pressure or temperature control,

unsuitable constraints and more. The simulation software needs to be carefully and properly configured as even the small change in parameters influences the simulation and can result in blowing up or in results that do not correspond with experimental results.

Limitations of classical MD are based on its approximations and on the characteristics of the method. MD with molecular mechanics model works at atomic level, the atom is an undivided, mass and charge point. Electrons are not explicitly accounted, their distribution is usually represented by one scalar, the partial charge. The topology of the system (how atoms are bonded) stays the same during the whole simulation and no chemical reactions can happen[1].

Potentials

All interactions are described by empirical functions that express the potential of interactions with respect to the distance between atoms (and other parameters). The form of these functions and the values of their parameters together compose a system called *force field*. Various force fields have been developed, they differ in suitability for particular application—some are more suitable for organic molecules, other focus on inorganic compounds. The most common are Amber [CCB+95], CHARMM [MBF01], MMFF94 [Hal96] and UFF [RCC+92].

All bonded interactions are short-ranged and they occur only between atoms connected by covalent bond. Bond stretching and angle bending are usually approximated by harmonic potential, torsion angles by cosine function. Short-range non-bonded interactions, such as van der Waals interactions, London dispersion interactions and repulsions due to Pauli exclusion principle, decay very fast with increasing distance. Therefore, they are usually calculated only between atoms within short distance. Most common approximation is Lennard-Jones potential [MSH07], expressed as

$$U_{Waals} = \sum_i \sum_{j>i} 4\epsilon_{ij} \left[\left(\frac{\sigma_{ij}}{\mathbf{r}_{ij}} \right)^{12} - \left(\frac{\sigma_{ij}}{\mathbf{r}_{ij}} \right)^{6} \right] \tag{6}$$

where ϵ_{ij} is the depth of a potential well, σ_{ij} is the distance at which the inter-particle potential is 0.

As mentioned, short-range interactions of the considered particle are calculated only with a limited number of particles that are closer than preset distance. A few methods can construct a list of neighboring particles for all atoms in $\mathcal{O}(N)$ time (for one atom in constant time) so that it is not necessary to check every pair of particles in the system each step [Pli95]. First approach, called Verlet lists, includes in the neighbor list of the considered particle all

[1] Although, force fields for some specialized cases of chemical reactions have been developed, e.g. [vDDLG01].

particles within extended cutoff distance $c_{extend} = c_{cutoff} + \delta$. Verlet list is rebuilt every few timesteps and δ is set so that no atom within c_{cutoff} can move further than c_{extend}. Second approach, called link-cell method, bins all atoms to cells and for the considered particle only particles in the same cell or neighboring 26 cells are examined. The more efficient combination of these methods uses binned cells to create Verlet list.

Short-range interactions, either bonded or non-bonded, are calculated very quickly as for each atom there is only small, upper-bound number of interacting atoms. They do not present a challenge from computational point of view but they crucially influence the movements of atoms.

2.2 Calculation of Long-Range Interactions

The bottleneck of the computation, electrostatic interactions ranging to long distances are calculated by Coulomb's law between almost every pair of atoms (only atoms with bonded interactions are excluded) as

$$U_{Coulomb} = \sum_i \sum_j \frac{q_i q_j}{|\mathbf{r}_i - \mathbf{r}_j|}. \tag{7}$$

The research in methods of molecular dynamics focused intensively to approximate the Coulomb's law and reduce the intrinsic quadratic complexity [Koe06]. The simplest method—the cutoff method—calculates the interactions by Coulomb law but only for atoms that are within preset cutoff distance; it completely neglects the interactions between atoms that are further apart. The method is rather inaccurate and it may introduce unphysical artifacts at the edge [LB89, SS92, BAD05], however, it is fast and easily implementable. Smoothed cutoff method and Wolf summation method [WKPE99] deals with the artifacts but the accuracy remains rather low.

All more sophisticated and more accurate methods interpolate the charges onto the grid and then they calculate the potential by

- Fourier transform: Ewald sum [Ewa21], Particle-Mesh Ewald [DYP93], Smooth Particle-Mesh Ewald [EPB+95], Particle-Particle Particle-Mesh method [HE88], Gaussian split Ewald [SKE+05];
- hierarchical division of the space: Barnes-Hut method [BH86], Fast Multipole Method [GR87];
- multigrid methods [SD01, STH02, SS05].

Fourier transform methods are based on Ewald sum [Ewa21]. The electrostatic interactions between close atoms are calculated precisely by Coulomb's law, the long-range contribution is interpolated onto the grid, calculated in Fourier space and then interpolated back from the grid. Most software packages for MD simulation implement them, along with periodic

boundary conditions they require. However, they do not scale well in number of processors used due to the remaining many-to-many communication pattern. Most of the methods based on Fourier transform have $\mathcal{O}(N \log_2 N)$ asymptotic complexity, however, they run equally fast or even faster than $\mathcal{O}(N)$ algorithms presented in the next paragraph. Implementation characteristics and code optimizations can efface the difference, moreover, $\log_2 N$ is below 30 even for the largest systems.

Methods based on the hierarchical division of space stem from astronomical N-body simulations that resemble MD (instead of electrostatic potential they calculate gravitational interactions). Fast multipole method hierarchically divides the simulation space into subcells and considers a cluster of particles as one particle with combined charge from faraway point. Electrostatic interactions between particles in same or neighboring cells (on the finest level) are computed directly, others are approximated by the multipole expansion. It has $\mathcal{O}(N)$ complexity, however with a large multiplicative constant. Moreover, the implementation is rather difficult.

Multigrid methods, a mathematical approach to solve partial differential equations [ST82, BvHM00], apply local process to multiple scales (grids) of the problem. The algorithm calculates in V-scheme of grids: it goes from the finest grid to the coarsest grid and then back. First, it obtains an initial approximation of results on the finest grid. As the error is smooth, the corrections can be calculated on coarser grids by recursive relaxation (e.g. Jacobi relaxation method) and restriction to coarser grids. The correction then prolongates from the coarsest grid to finer and finer grids. Their large multiplicative constant for $\mathcal{O}(N)$ complexity and iterative nature makes them difficult to implement for massively parallel resources.

Multilevel summation method [Har06] with $\mathcal{O}(N)$ complexity resembles multigrid methods. It calculates differently varying parts of potential on grids with different spacing. It exhibits valuable advantages: one-to-many communication pattern, rather simple parallel implementation and reasonable number of floating point operations. Therefore, we have selected it as the fine function $\mathcal{F}$ for the parareal method. Nevertheless, there are no characteristics of $\mathcal{F}$ function preventing other methods to take its place.

2.3 Multilevel Summation Method

Hardy in his dissertation [Har06] developed the multilevel summation method (MSM) for calculation of long-range interactions with thorough mathematical background, performance assessments, accuracy analysis and implementation suggestions. The method divides the calculation of the potential onto multiple grids. To keep reasonable accuracy, more and more slowly varying parts of the potential are calculated on coarser and coarser grids.

Multilevel summation method hierarchically interpolates smoothed po-

tential onto multiple grids and then sums it up [HSS09]. The method therefore depends on two main functions—the smoothing function and the interpolation function.

With the smoothing function $g_a(\mathbf{r}_i, \mathbf{r}_j)$, the reciprocal distance $\frac{1}{|\mathbf{r}_j - \mathbf{r}_i|}$ can be rewritten so that it gives the same result but its parts correspond to more and more slowly varying parts of potential.

$$\begin{aligned} \frac{1}{|\mathbf{r}_j - \mathbf{r}_i|} &= \left(\frac{1}{|\mathbf{r}_j - \mathbf{r}_i|} - g_a(\mathbf{r}_i, \mathbf{r}_j) \right) + g_a(\mathbf{r}_i, \mathbf{r}_j) \\ g_a\left(\mathbf{r}_i, \mathbf{r}_j\right) &= \frac{1}{a}\gamma\left(\frac{|\mathbf{r}_j - \mathbf{r}_i|}{a}\right) \\ g_a\left(\mathbf{r}_i, \mathbf{r}_j\right) &= \left(g_a(\mathbf{r}_i, \mathbf{r}_j) - g_{2a}(\mathbf{r}_i, \mathbf{r}_j)\right) + g_{2a}(\mathbf{r}_i, \mathbf{r}_j) \end{aligned} \tag{8}$$

where a is the cutoff distance and γ is the smoothing function.

The method has multiple grids in a sequence and every grid considers doubled cutoff distance than the grid before. So, the grid in level k has the cutoff distance $2^k a$, where a is the cutoff on the finest grid. Equations (9) reformulates equations (8), instead of subscript a representing the cutoff distance, they use superscript k representing the grid level. We will use the superscript notation from now on.

$$\begin{aligned} \frac{1}{|\mathbf{r}_j - \mathbf{r}_i|} &= (g^* + g^0 + g^1 + ... + g^{l-2} + g^{l-1})(\mathbf{r}_i, \mathbf{r}_j) \\ g^*(\mathbf{r}_i, \mathbf{r}_j) &= \frac{1}{|\mathbf{r}_j - \mathbf{r}_i|} - \frac{1}{a}\gamma\left(\frac{|\mathbf{r}_j - \mathbf{r}_i|}{a}\right) \\ g^k(\mathbf{r}_i, \mathbf{r}_j) &= \frac{1}{2^k a}\gamma\left(\frac{|\mathbf{r}_j - \mathbf{r}_i|}{2^k a}\right) - \frac{1}{2^{k+1} a}\gamma\left(\frac{|\mathbf{r}_j - \mathbf{r}_i|}{2^{k+1} a}\right) \text{ for } k = 0..l-2 \\ g^{l-1}(\mathbf{r}_i, \mathbf{r}_j) &= \frac{1}{2^{l-1} a}\gamma\left(\frac{|\mathbf{r}_j - \mathbf{r}_i|}{2^{l-1} a}\right) \end{aligned} \tag{9}$$

where l is number of grids (grid levels $0..l-1$).

γ is an unparameterized smoothing of the function $\frac{1}{\rho}$ chosen so that the first part of the first equation (8), $\left(\frac{1}{|\mathbf{r}_j - \mathbf{r}_i|} - g_a(\mathbf{r}_i, \mathbf{r}_j)\right)$, vanishes after the cutoff distance (a for the finest grid, $2^k a$ for grids of level k). In the second part, $g_a(\mathbf{r}_i, \mathbf{r}_j)$, the function γ ensures partial derivatives and slow variation. Hardy usually sets it as Taylor smoothings, e.g.

$$\gamma(\rho) = \begin{cases} \frac{15}{8} - \frac{5}{4}\rho^2 + \frac{3}{8}\rho^4, & \rho < 1, \\ 1/\rho, & \rho \geq 1 \end{cases} \tag{10}$$

The interpolation function $\mathcal{I}$ makes it possible to do the following ap-

proximation

$$\begin{aligned}\frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} &= (g^* + g^0 + g^1 + ... + g^{l-2} + g^{l-1})(\mathbf{r}_i, \mathbf{r}_j) \\ \approx (g^* + \mathcal{I}^0(g^0 + \mathcal{I}^1(g^1 + ... + \mathcal{I}^{l-2}(g^{l-2} + \mathcal{I}^{l-1}g^{l-1})...)))(\mathbf{r}_i, \mathbf{r}_j)\end{aligned} \tag{11}$$

$\mathcal{I}$ interpolated g to the grid by

$$\begin{aligned}\mathcal{I}^k g(\mathbf{r}_i, \mathbf{r}_j) = \sum_\mu \sum_\nu \phi^k_\mu(\mathbf{r}_i) g(\mathbf{r}^k_\mu, \mathbf{r}^k_\nu) \phi^k_\nu(\mathbf{r}_j) \\ k = 0, 1...l-1\end{aligned} \tag{12}$$

where ϕ^k_μ is the nodal basis function with local support (the function is non-zero only in the close surrounding of the grid point). For each grid with spacing $2^k h$, where h is the spacing of the finest grid, and grid points r^k_μ, nodal basis functions ϕ^k_μ are defined as

$$\phi^k_\mu = \Phi\left(\frac{x - x^k_\mu}{2^k h}\right) \Phi\left(\frac{y - y^k_\mu}{2^k h}\right) \Phi\left(\frac{z - z^k_\mu}{2^k h}\right), \tag{13}$$

where $\Phi(\epsilon)$ can be for example the cubic interpolating polynomial[2].

With approximated reciprocal distance, the potential is calculated as

$$U_i \approx \frac{1}{4\pi\epsilon_0}(u^{short}_i + u^{long}_i) \tag{14}$$

The following pseudoalgorithm represents the equation written above. The short-range part within cutoff is $u^{short}_i = \sum_j g^*(\mathbf{r}_i, \mathbf{r}_j) q_j$. Long-range part is recursively divided between two parts—one within the cutoff and then calculated on the current grid; and the second representing even more slowly varying potential and then calculated on coarser grids. The method for calculating the long-range part that corresponds to evaluation of equation (11) goes as follows [HSS09]:

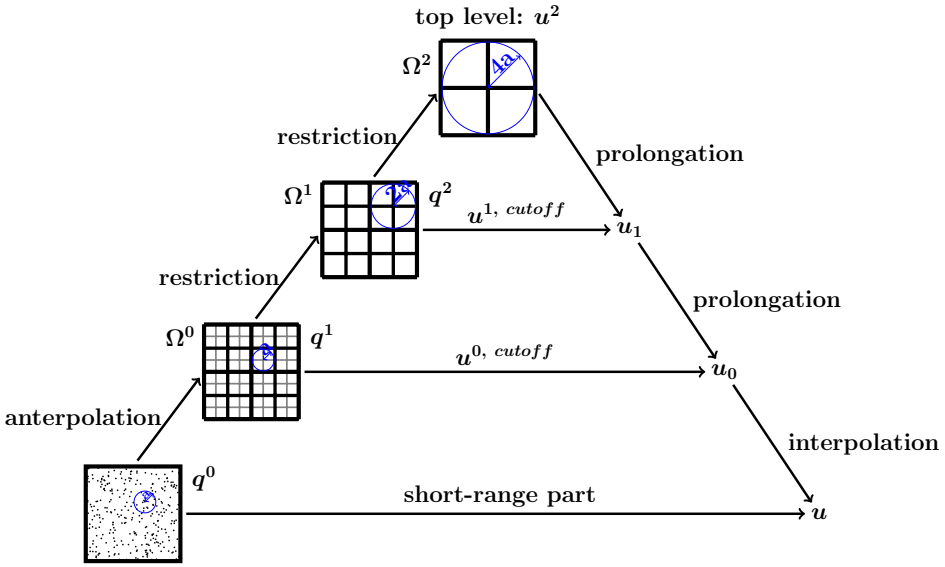
- **anterpolation**—puts point charges onto the grid $q^0_\mu = \sum_j \phi^0_\mu(\mathbf{r}_j) q_j$;
- recursively for $k = 0, 1, ..., l-2$
 - **restriction**—approximates charges onto coarser grid $q^{k+1}_\mu = \sum_\nu \phi^{k+1}_\mu(\mathbf{r}^k_\nu) q^k_\nu$;
 - **lattice cutoff**—calculates the part of potential corresponding to the grid $u^{k,cutoff}_\mu = \sum_\nu g^k(\mathbf{r}^k_\mu, \mathbf{r}^k_\nu) q^k_\nu$;

[2]Letters μ and ν represent grids or grid points, k is the grid level, h is the spacing of the finest grid.

- **top level**—calculates the most slowly varying part of the potential corresponding to the coarsest grid $u_\mu^{l-1} = \sum_\nu g^{l-1}(\mathbf{r}_\mu^{l-1}, \mathbf{r}_\nu^{l-1}) q_\nu^{l-1}$;
- **prolongation**—recursively backwards for $k = l-2, ..., 1, 0$ adds up the parts of potential corresponding to the grids $u_\mu^k = u_\mu^{k,cutoff} + \sum_\nu \phi_\nu^{k+1}(\mathbf{r}_\mu^k) u_\nu^{k+1}$;
- **interpolation**—puts grid potential off grid $u_i^{long} = \sum_\mu \phi_\mu^0(\mathbf{r}_i) u_\mu^0$.

The Figure 2.1 shows a specific example with three grids.

Figure 2.1: Example of multilevel summation algorithm.



Hardy precisely calculated the number of floating point operations needed for the evaluation of the multilevel method [Har06, ch. 2]. It is

$$\left(\frac{4}{3}\pi m + \frac{32}{3}\pi + \frac{81}{2}\right)\left(\frac{a}{h^*}\right)^3 N + \left(6p^3 + 31p^2 + 36p + 17\right) N + \left(\left(\frac{4a}{h}\right)^3 + 14(p+2)\right)\frac{8}{7}\left(\frac{h^*}{h}\right)^3 N \tag{15}$$

where

- $h^* = N^{-1/3} L$ is the average distance between nearest neighbors
- h is the spacing in the finest grid
- a is the cutoff distance

- N is the number of atoms
- L is the length of the box
- γ is the smoothing function polynomial of degree $2m$
- ϕ is the basis function polynomial of degree p

For common parameters; $h^* = 1\,\text{Å}$, $h = 2\,\text{Å}$, $a = 12\,\text{Å}$, $m = 2$ (for C^2 Taylor smoothing function), $p = 3$ (Hermite cubic interpolant) with relative error in force less than 1% [Har06, p. 78]; the number of floating point operations reaches

$$77.7a^3N + 566N + \frac{73a^3}{h^6}N + \frac{80}{h^3}N \approx 136813N. \tag{16}$$

Multilevel summation method exhibits several advantages over common methods for calculation of electrostatic interactions [Har06, ch. 1.2]. In comparison with fast multipole method (FMM), MSM calculates continuous forces and smooth potential. Therefore lower accuracy (e.g. lower order of interpolants) suffices for stable dynamics and conservation of energy. MSM, as multigrid and Fourier based methods, does not conserve linear momentum, FMM does. It is expected to fluctuate but not drift [STH02], so it does not present an issue. When compared to Particle-Mesh Ewald method (PME), multilevel summation method scales better in parallel computation, uses multiple-time-step integration scheme more efficiently and communicates less [Har06, ch. 7.4]. Apart from multigrid methods, MSM is not iterative therefore it scales better in parallel computation and does not require more than one global communication exchange.

2.4 Acceleration of Molecular Dynamics Simulations

Common MD experiments simulate hundreds of nanoseconds. However, many interesting biological and chemical processes occur at longer timescales: tens of microseconds and more. Due to 2 fs timestep, these long simulations require billions of steps, each computationally demanding. This subsection presents techniques that enable current MD software reach simulation times longer than hundreds of nanoseconds through model simplifications, computational volume reduction, hardware acceleration and parallel calculation.

The first type of methods further approximates and simplifies the atomic step-by-step model of MD simulations with molecular mechanics. Higher abstraction leads to less demanding computation of each step or to less steps needed to simulate whole process of interest.

Coarse models [RB98] describe interactions of beads that consist of a few (usually tens of) atoms. The timestep of MD can be increased as the fast

movements of classical MM MD are hidden within the bead. Thanks to that, molecular dynamics with coarse-grained models simulates up to 1 ms processes.

Discrete molecular dynamics [PDD11] does not integrate Newton equations after fixed timestep, it assumes the constant velocity of every atom until a collision occurs. Then it recalculates velocities of collided atoms so that the momentum is conserved. As the interactions are calculated only within small number of atoms (only within those taking part in collision) and usually less frequent than every 2 fs, the simulation goes significantly faster compared to the classical molecular dynamics.

Conformation space mapping and energy surface approximation can complete their task much faster when the simulation introduces biased, artificial potential to the system to avoid staying in local energy minima. Metadynamics [LG08], umbrella sampling [TV77], accelerated molecular dynamics [MF04], replica exchange [SO99] and similar address the sampling problem with good results.

Also, implicit models of solvent described below belong to this group.

The second type of methods eliminate the parts of computation that usually give almost the same values, such as interactions within the water molecule.

Most of the atoms in many simulated systems construct water molecules: up to 90%. Systems need to be properly solvated due to two reasons. First, they give results better corresponding with reality. Without explicit water molecules, systems often include non-realistic physical artifacts due to changed electrostatics. Second, periodic boundary conditions, common in MD simulations, need at least 10 Å layer of water to ensure that two replicas of the molecule of interest do not "see" each other—their interactions are close to 0, thus negligible. In explicit modelling, water usually has a rigid model that saves the calculation of interactions within the water molecule [HSP+04, Ric04, TR85]. Implicit modelling does not explicitly consider solvent molecules, it includes its effects into the parameters of the system [VM05, Bak05]. That greatly reduces the number of atoms, shortens the time to result and enables quite fast multimillion-atom simulations such as [SLPS09]. Hybrid models [LSOS04] combine these two approaches: at the boundary of the molecule of interest, explicit water molecules interact with it, further away there is the reaction field. However, at the interface between explicit and implicit representation, explicit water molecules do not behave in correspondence with reality.

Thanks to constraints of bonds that contain the hydrogen, the MD simulations have 2 fs timestep. These bonds vibrate at frequency $10^{14}\ s^{-1}$ so the timestep should be 1 fs, constraints double that number. Most common algorithms are SHAKE [RCB77], LINCS [HBBF97, Hes08], SETTLE [MK92], RATTLE [And83]. Force fields with united atom model (for example Amber) treat hydrogens differently than other atoms, they "merge" them with

bonded atom (usually oxygen, carbon or nitrogen) and alter the parameters for the united atom [PCC+95].

Multiple-time-step integration scheme takes into account that different interactions fluctuate over different timescales [AT89, GT98]. For example, the potential due to bond stretching changes its value more often than the van der Waals potential or electrostatic potential. Therefore electrostatic interactions can be evaluated less frequently than the bonded interactions.

The third type of approaches relies on hardware. One of the longest simulations has been conducted on specialized hardware Anton in the research of Shaw et al. [SDD+07] calculating ten microseconds of simulation time per day. The machine computation time has been offered for research calculations for free but naturally it is not capable to satisfy all needs of the whole scientific community. Its cost prevents this concept to spread widely. Several MD simulation algorithms have been implemented on GPU [SPF+07, ALT08, LSVMW08, HSS09] with speed-up up from 10 to 100 (with multi-GPU even by three orders of magnitude) compared to optimized CPU code.

Parallel and distributed implementations of molecular dynamics code naturally speed up the calculation; this topic is analyzed in the next subsection *Parallel Computation and Scalability*.

Presented methods and improvements make it possible to simulate even large systems (millions of atoms) for quite long simulation time (hundreds of nanoseconds) on rather available resources (thousands of cores). The acceleration leads to two goals: it is possible to simulate larger systems or it is possible to simulate the same systems faster. If the same simulation takes less wallclock time, longer simulation timescales are feasible and researchers from biology, chemistry or medicine can observe more and more interesting processes at atomic scale.

Parallel Computation and Scalability

The acceleration and parallel runs of MD algorithms follow two goals. First, to cut the time to result. The wallclock time of the computations crucially influences their usability and impact. The simulation that takes tens of minutes to complete on available resources can be repeated several times a day and the researcher can faster reveal mistakes or verify a hypothesis—the simulations push the research further. With short time to result, longer simulation times become feasible and previously unsimulated phenomena can be studied through MD simulations.

Second, to enable simulations with larger number of atoms. Fortunately, the simulations with more atoms but also more computational resources tend to keep the wallclock time similar. However, interesting changes in larger systems take longer, which leads again to the necessity of longer simulation times.

Types of scaling examined in MD algorithms correspond with the two goals. First, the strong scalability examines how the wallclock time of fixed-size simulation changes (ideally reduces) as the number of computational resources grows. For example, an algorithm with an ideal strong scaling would simulate the same system million times faster on a million-core supercomputer than on a one-core computer. Naturally, every algorithm eventually hits the strong scalability wall, when the fine granularity of work reaches the critical level and the wallclock time increases due to the communication and synchronization overhead. The research of parallel/distributed implementations of MD algorithms wants to shift the wall further away and take advantage of massively parallel infrastructures.

Second, the weak scalability examines how the wallclock time changes (ideally remains the same) as both the size of the problem (number of atoms) and the number of computational resources grow, i.e. the amount of work per processor remains the same. For example, an algorithm with an ideal weak scaling would simulate 2000-atom system on two cores for the same wallclock time as 1000-atom system on one core. Weak scalability makes it possible to simulate large system without major problems if enough computational power is provided. However, as mentioned above, large size of the systems goes in hand with longer timescales necessary for relevant changes to happen.

All common MD algorithms have their parallel implementations with the spatial decomposition in common software packages such as Gromacs [HKvdSL08], Amber [CCD+05], NAMD [PBW+05], GROMOS [CHB+05], LAMMPS [Pli95, Pli03] or in libraries such as Scafacos [ABF+13, BFH+]. Bowers et al. in [BDS05] overviews usual decomposition techniques. All parallel implementations exhibit almost perfect weak scaling, the evaluations of strong scalability show the wall commonly at tens of thousands of processors. The algorithms with the highest strong scalability, to the author's best knowledge, achieve the saturation up to half a million cores.

Andoh et al. [AYF+13] conducted a simulation with 10^7 atoms on over half a million cores with simulation speed 35 ns/day, one step was evaluated in 5 ms. The long-range interactions were calculated by the Fast Multipole Method. Richards et al. [RGC+09] achieved almost 260 TFLOP/s performance in the simulation with $\sim 10^9$ particles and showed almost ideal strong scaling up to 294000 cores. Copernicus [PLP+11] has simulated a folding process of a villin protein (10^5 atoms) through a series of short simulations and Markov chain that connected them. They managed to reach remarkably fine granularity, 5000 cores simulated only 10000-atom system in wallclock time 3.5 day. Their results show the potential of time parallelism incorporated into MD simulations.

The spatial decomposition does not suffice in three cases [SC05]. First, when the simulated process has small number of atoms but it takes rather long. Second, when the amount of work per processor gets low and the fine

granularity stops the acceleration. And third, when the communication is expensive as is common in distributed environments. All these three cases are becoming more and more common. Researchers want to simulate long processes due to their chemical or biological relevance but algorithms struggle to efficiently utilize increasing number of (sometimes distributed) computational resources at hand to make that possible.

Therefore, new approaches are needed to saturate more computational power and use it to shorten the time to result. All state-of-the-art MD algorithms decompose only spatial domain and run parallel-in-space. We believe that the parallelization along the time domain would increase the level of parallelism and shift the strong scalability wall further.

2.5 Parallel-in-Time Computation

In parallel-in-time computation, different processors calculate results in different time points. Traditional, parallel-in-space computation decomposes the spatial domain; different processors calculates results for different points in space. However, decomposition of the temporal domain is not as easy due to the sequential characteristic of time. Molecular dynamics sequentially solves the initial value problem, so the result from previous time points determines the result in the next time point. Several techniques are able to overcome this requirement, we introduce those concerning MD.

The project Copernicus [PLP+11] and Folding@Home [LSSP02] use coarse-grained time parallelism. Many simulated processes include long but uninteresting metastable states of local minimum in which the system gets stuck until it crosses the energy barrier. For example, when a protein is folding, it goes through different conformations, each of them is in the local minimum of energy surface. The interesting part—how the protein gets from one conformation to another—is rather quick compared to the stay in metastable state. Projects build the Markov model of different conformations that gradually explores the conformation space by many short simulations. That allows them to use highly distributed framework and achieve remarkable strong scaling. Two different processors can simulate two possible paths from one conformation at the same time, hence the time parallelism. This approach is suitable for simulation of processes that consist of many metastable states separated by short transitions, e.g. protein folding or conformation space exploration.

Yu et al. in [YSC06] use data from prior related simulations to guide the system and predict its changes. The simulation time is divided into intervals given to different processors. The prediction algorithm forecasts the state of the system at the beginning of each interval according to previous simulations. That gives the initial value for every processor and the classical MD simulation through the interval follows. At the end of the interval the processor verifies that the state corresponds with the pre-

dicted state at the beginning of the next interval. If not so, the data from the mispredicted point in time are discarded, prediction algorithm adapts and the simulation continues from the latest correct state. The method has been evaluated on nanotubes pulled by the external force, they achieved high strong scaling and speed-up. However, the method relies heavily on reliability of the prediction: the presented high speed-up and strong scalability depend on no mispredictions during the simulation. The prediction algorithm for nanotubes has been developed, no other suggestions for other MD applications (trajectory examination, conformation space exploration) have been mentioned.

Several mathematical methods have been developed that calculate parallel-in-time in fine granularity: they compute results of a time dependent differential equation in a few (successive) time points simultaneously [GV07]. The first such method by Nievergelt, 1964, [Nie64] later became known as multiple shooting method. It divides the time intervals to many subintervals, then solves initial value problem for each subinterval and forces continuity by Newton procedure. Time-parallel approach to iterative methods for solving partial differential equations with implicit integration schemes [DMDS95] were followed by applying multigrid methods for acceleration [Hor92, VvdV94, HV95]. In 2001, Lions, Maday and Turinici introduced *parareal in time* method [LMT01]that has been extensively analyzed since then [GV07, Mad08, MT05, Aub11].

Baffico et al. in [BBM+02] have done the first MD simulation with the parareal algorithm. The rather limited paper examined the possible speed-ps and suitability of the parareal scheme for classical MD and *ab initio* MD that considers quantum mechanics as model for interactions. Shorter and longer timesteps were applied for the fine and coarse function, respectively. They concluded that this approach is worth exploring and with many possible choices of fine and coarse functions, we are "limited only by our imagination".

Waisman and Fish [WF06] combined the multigrid method [SD01, STH02, SS05] and waveform relaxation method [LRSV82] into space-time multilevel method with implicit integration scheme and speeded up MD simulation of a polymer melt.

In 2013, Bulin published a master thesis at Stockholm university [Bul13] where he compared the waveform relaxation method [LRSV82, VvdV94] and the parareal method [LMT01]. Moreover, he suggested a few improvements that could increase the scalability of these methods. For the parareal method, he proposed an intuitive windowing (explained in the following subsection) and applied multiple levels into the parareal scheme. He achieved speed-up up to 10, the number of iteration ranged from 4 up to 25 for different widths of the window. Bulin concluded that the waveform relaxation function is "useless for this kind of problems" due to slow convergence. For the parareal algorithm, he stated that it is not suitable for large

scale computing due to low speed-up. However, different coarse functions (faster, yet still reasonably accurate and numerically stable) could overcome this.

Despite rather disappointing results of parareal scheme in MD simulations, we want to research it further. Speck et al. in [SRK+12] managed to simulate large gravitational N-body system through an altered parareal method that nicely combined the parallelization in spatial and temporal domain. As the electrostatic and gravitational N-body problems have much in common, the parareal scheme with a more appropriate coarse function may succeed in MD simulation.

2.6 Parareal Method

The parareal method [LMT01, BM02, Mad08] parallelizes the time domain by approximating the solution in time t_{i+1} without accurate solution in time t_i where $t_{i+1} > t_i$. When solving differential equation, we seek the function of time $\mathbf{u}$ for initial condition v in time $t > 0$. The exact solution is rarely known, more often we have the precise enough approximation $\mathcal{F}_\tau(t, v)$ obtained by discretization with small timestep. $\mathcal{F}$ approximates $\{\mathbf{u}(T_n)\}_n$ as

$$\{\lambda_n = \mathcal{F}_{T_n - T_0}(T_0; v) = \mathcal{F}_{\Delta T_n}(T_{n-1}; \lambda_{n-1})\}_n \tag{17}$$

where T_i are evenly-spaced time points. The sequential character of the problem appears clearly: the solution in time T_n can be calculated only after the solution in time T_{n-1} is known.

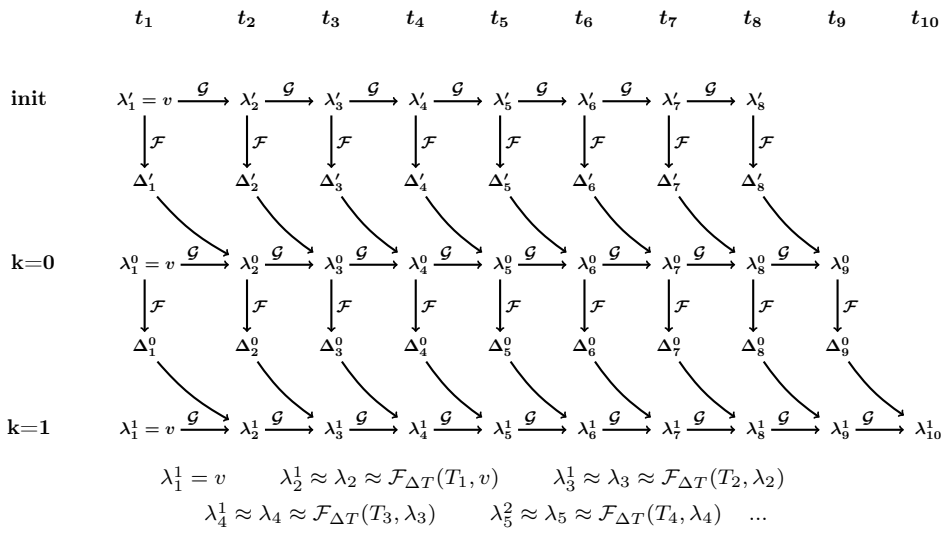
The parareal method proposes a sequence $\{\lambda_n^k\}_n$ that converges to $\{\lambda_n\}_n$ rapidly as k, the number of iteration, increases and that can be built in parallel. It introduces the second, coarse and cheap approximation $\mathcal{G}_\tau(t, v)$. The sequence $\{\lambda_n^k\}_n$ is then defined recursively as

$$\lambda_{n+1}^{k+1} = \mathcal{G}_{\Delta T}(T_n; \lambda_n^{k+1}) + \mathcal{F}_{\Delta T}(T_n; \lambda_n^k) - \mathcal{G}_{\Delta T}(T_n; \lambda_n^k) \tag{18}$$

The idea behind this sequence is to shift the inherent sequential nature of calculation from $\mathcal{F}$ to $\mathcal{G}$. The equation (18) resembles predictor-corrector integration scheme [ML67, Gea71]. Function $\mathcal{G}$ roughly assesses the initial approximation of the results. The difference between the results from precise calculation and from coarse calculation on the same data presents the error that is included into calculation in the next iteration. That gradually improves the approximation of the result.

In a long simulation with the parareal scheme, the width of "computational window" is determined by properties of the fine and the coarse function. Within the computational window, the results in successive time points are calculated parallel-in-time. As the calculation proceeds (as explained in the next paragraph), results from time points in the computational window converge and the window shifts to the right on the time axis.

Figure 2.2: Computational flow of the parareal method.



The Figure 2.2 shows an example of the parareal method that calculates λ_n for ten time points $n = 1..10$ with the initial condition v. The results converge after two iterations $k = 0, 1$. We apply further notation for clarity:

- $f_n^k = \mathcal{F}_{\Delta T}(T_{n-1}, \lambda_{n-1}^k)$
- $g_n^k = \mathcal{G}_{\Delta T}(T_{n-1}, \lambda_{n-1}^k)$
- $\Delta_n^k = f_n^k - g_n^k$
- $'$ superscript replaces $(-1)^{st}$ iteration—initialization

The arrows in the Figure 2.2 show the computational flow—what needs to be computed in what order and with what dependencies. In *init* row, the calculation of g'_n depends on already known g'_{n-1}. Downward arrows with f_n^k and Δ_n^k represent parallel computation of the error between the precise approximation and the coarse approximation done on the same data. In k rows, λ_n^k calculation depends on known g_{n-1}^k and Δ_{n-1}^{k-1}. After two iterations, we consider λ_n^1 very close to the result (e.g. $\Delta_n^k < \epsilon$) of the sequential computation of precise approximation based on accurate input data. The converged result in the last time point, $lambda_{10}^1$, will serve as the initial condition for calculation in the next computational window.

The parareal algorithm multiplies the volume of computation, i.e. the number of floating point operations needed to get the results. Moreover, it requires several times more computational resources than the traditional, sequential-in-time integration schemes. However, if the number of iterations is smaller than number of time points, the method shortens time to result.

2.7 Limitations of the Parareal Method

The parareal method has its limitations: the most important is the relation between convergence and scaling. If the method calculates in too many time points simultaneously, the convergence begins to fail and the number of iterations increases. And also too inaccurate coarse function with poor convergence and high number of necessary iterations assess well only results within a few time points ahead.

The maximum theoretical speed-up of the method depends mainly on the ratio of evaluation time of the expensive $\mathcal{F}$ function and the cheap $\mathcal{G}$ function, high ratio results in high theoretical speed-up. The high ratio is ensured by low evaluation time of $\mathcal{G}$ function, so cheaper functions are preferred. However, the cost of the coarse function coincides with accuracy. Cheap but too inaccurate functions cause lower convergence and higher number of iterations.

MD Simulations so far conducted with the parareal method [BBM+02, SC05, NVK93, Bul13, BLMM13] achieved only low speed-up and efficiency or they exhibit issues with convergence. All of them used $\mathcal{G}$ function based on larger timestep. We believe that such function is unsuitable as it cause numerical instability in molecular dynamics and with more appropriate functions we can achieve better results.

The convergence of the parareal method has been extensively researched for both ordinary and partial differential equations, usually for Euler integration schemes [MT05, Mad08, GV07, Bal05, SRn03]. It has been proven that with any function $\mathcal{G}$, with no requirements on its quality, the parareal scheme will converge after at most $T - 1$ iterations, where T is the number of time points [GV07]. Of course, in order to achieve some speedup, the number of iterations has to be much lower than number of time points, i.e. $K \ll T$. That depends on the stability of $\mathcal{G}$ function.

2.8 Summary

Molecular dynamics offers biologists and chemists the opportunity to observe processes at high resolution both in space and time. This work deals with molecular dynamics with molecular mechanics model of interactions. MM MD sequentially calculates the forces exerted on all atoms caused by their interactions and then moves them according to Newtonian physics. The potentials described by empirical functions are evaluated in every step, they are computationally demanding due to the long-range electrostatic interactions. Moreover, due to small timestep (usually 2 fs), it takes large number of steps to reach interesting timescales (more than 10^6 for short, 2 ns simulations). Many approximations for Coulomb potential have been developed and their parallel implementations along with other improvements have made it possible to simulate millions of atoms for hundreds of nanosec-

onds. However, with increasing number of computational resources, issues with strong scalability have arisen. Only a few of current algorithms are able to saturate over 10^5 cores and their ability to run faster when provided with more processors hits the limit. As all of them decompose only along the spatial domain, we believe that computation parallel both in time and space would push the level of parallelism further, thus increase the strong scalability and harness even massively parallel computing infrastructures. First attempts to calculate MD simulations parallel-in-time have been made, however with little success. We want to study and develop the parareal method and incorporate it in MD code. Apart from published approaches, we have selected different, more appropriate coarse functions that should provide high speed-up at relatively low cost and reasonable convergence.

3 Research Questions and Proposed Solutions

3.1 Research Questions

MD simulations have proven their usability for chemists and biologists. With faster algorithms and more computational resources, the researchers naturally wish to simulate larger systems for longer simulation times. Parallel implementations of current methods manage simulations with higher number of atoms and short simulation time without major problems [3]. However, processes concerning large systems usually take more time: it would be useless from biological point of view to simulate a whole cell with 10^{14} atoms for a nanosecond. Moreover, even many small systems take part in rather long processes of great biological or chemical interest. Current algorithms have almost reached their potential to cut the time to result by adding more computational power that is available.

A virus passing a cell's membrane is an example of rather long process (at least 50 μs) with large number of atoms (hundreds of millions when properly solvated). Vácha et al. has performed coarse-grained simulations in implicit solvent [VMVF11], however they cannot answer some questions of chemical interest such as whether the water molecules surround the virus even when it is enveloped in the membrane after the passing. The full-atom simulation of properly solvated system would help them to observe the process at close look. As computer scientists, we see the challenge in *how we can simulate such a large system for such a long simulation time*.

One possible way is to reduce the number of floating point operations by changing the MD algorithm, e.g. by inventing the algorithm for calculation of long-range interactions with complexity $\mathcal{O}(\log N)$. However, that difficult task requires thorough background in chemistry, physics (to assess correctly what approximations can be made without major influence on accuracy), mathematics and computer science (to formulate N-body problem in such way that an efficient implementation would be possible).

Another approach is to increase the level of parallelization and utilize the increasing number of computational resources at hand. Spatial decomposition that lays behind all current parallel implementations of MD code does not suffice for long simulation times as it scales only with the size of the system, not the number of steps. The strong scalability of current algorithms hits the wall at hundreds of thousands of processors. The largest supercomputer today has over three million cores[4] and prognoses suggest that exascale supercomputer will be built till 2020 with $\sim 10^9$ cores. Therefore,

[3] The issues connected with large simulations cannot be underestimated. However, they are more of a technical character as opposed to problems with long simulations times that come from the fundamental concept of MD.

[4] According to November 2013 list at http://www.top500.org, top 1 is Tianhe-2 at National Super Computer Center in Guagzhou, China, with 3 120 000 cores and almost 55 TFLOP/s peak performance.

we want to push the strong scalability further to use more computational resources to shorten the time to results. We strongly believe that parallelization in time would do that as the scaling will then depend also on the number of steps not only on the number of atoms.

Time parallelism has been discovered for MD simulations. Coarse-grained parallelism in time has been developed in Copernicus and Folding@Home projects [PLP+11, LSSP02]. Even though their contribution to the computational chemistry is undoubtful, they can not be applied for simulation of every biological process. The presence of metastable states and short length of interesting parts are two main conditions that not every process can satisfy. For example, when the virus passes through the cell's membrane, it is a slow but continuous process researchers want to observe step-by-step. Yu et al. [YSC06] guide the simulations of nanotube with the previously acquired data from similar *in sillico* experiments. However, it heavily relies on the prediction algorithm. We consider to apply their approach in some way also to our method. Fine-grained time parallelism represented by mathematical parallel-in-time methods combined with MD have rather disappointed in previous research. Especially, the parareal method has been analyzed, but the coarse function $\mathcal{G}$ based on larger timestep led to unsatisfactory convergence and speed-up.

We believe that with coarse functions based on more appropriate concepts, we can achieve better results. As the fine function $\mathcal{F}$, we have chosen the multilevel summation method. Therefore, we propose the parareal MSM method, details follow in the next subsection.

3.2 Proposed Solutions

Our main research question focuses on how to simulate molecular dynamics of large systems for long simulation times. We want to increase the strong scaling and the level of parallelism by parallel-in-time algorithm to achieve computation parallel in both space and time. We propose a novel method that combines the multilevel summation method and parareal time integration method into parareal multilevel summation method: MSM will calculate parallel-in-space; the parareal method will calculate parallel-in-time. The fine function $\mathcal{F}$ in parareal scheme will be classical MD with the multilevel summation method for calculation of long-range interactions. The main characteristic of the combination is the choice of parareal method's coarse approximation $\mathcal{G}$. This function can be based on various concepts:

- further simplification of the model: discrete MD;
- different parameters of MD algorithm: longer timestep;
- different parameters of the method for evaluation of long-range interactions: coarser grid in MSM, shorter cutoff in MSM;

- less expensive method for evaluation of long-range interactions: simple cutoff method, Wolf summation method.

We have considered several approaches and compared them to classical molecular dynamics with MSM. Moreover, we have examined two key aspects—the convergence and the ability to run in parallel.

Discrete Molecular Dynamics Coarse approximation $\mathcal{G}$ would run the simulation using another concept—instead of integrating the differential equations to capture the movements of atoms we assume that they move with constant velocity unless a collision occurs [PDD11]. Despite its very good convergence, discrete molecular dynamics would not work as it produces completely different trajectories than classical molecular dynamics. The overall properties of the system are similar but the atoms move differently. As we are particularly interested how exactly atoms move (for example, as the virus passes the membrane), this approach does not suit us.

Longer Time Step Coarse approximation $\mathcal{G}$ would have longer timestep in integration scheme, instead of 2 fs it would be e.g. 10 fs. The problem would appear in the beginning of the computation, before $k = 0$, when we need to calculate $\lambda'_{n+1} = \mathcal{G}(t_{n+1}, t_n, \lambda'_n)$ where $\lambda'_1 = v$. The coarse approximation $\mathcal{G}$ therefore needs to give at least a little reasonable results for every t_n knowing only initial condition v and its own result for t_{n-1}. Unfortunately, MD simulations with timestep larger or equal 5 fs give unusable results after just a few steps (in our own experiment with retinol and timestep 10 fs the system blew up after two steps). Half-converged results, e.g. λ^3_n, instead of λ'_n, could save the convergence but we would have to shorten the computational window. That would drastically reduce the ability to run in parallel. A few simulations of the parareal scheme with longer timestep coarse function have been conducted [BBM+02, Bul13, BLMM13], with relatively modest speedup for larger simulations (lower than 10) and issues with convergence (up to 25 iterations).

Coarse Grid in MSM Coarse approximation $\mathcal{G}$ would run molecular dynamics with multilevel summation method for evaluation of long-range interactions but the finest grid would have spacing e.g. four times larger than in the fine approximation $\mathcal{F}$. It would be easily implementable, however, the complexity assessment suggests it would not offer high speed-up as the number of floating point operations depends mainly on the cutoff distance a, not on the finest grid spacing h.

Shorter Cutoff Distance in MSM Coarse approximation $\mathcal{G}$ would run molecular dynamics with MSM but the interactions on the current grid will

be calculated within shorter cutoff range[5] than in the fine approximation $\mathcal{F}$ (e.g. 6 Å vs 12 Å for the finest grid). Larger parts of the potential would be calculated on coarser grids, eventually large part of the potential would be calculated on the top level where it is cheaper than on finer grids. In evaluations [Har06, ch. 3-5], Hardy found the relative error of forces to be in range from 1% to 5% with cutoff beginning at 8 Å. Our analysis showed that the maximal theoretical speed-up is 7. As simple cutoff method or Wolf summation offer much larger theoretical speed-ups, we abandon this possibility.

Cutoff Method Coarse approximation $\mathcal{G}$ would calculate long-range interactions by the simplest method available—sum up the interactions with Coulomb equation between atoms within constant cutoff, usually set to 12 Å. The accuracy does not reach the level of elaborate methods but the precise $\mathcal{F}$ function should correct the errors. The rather good convergence of the method should enable large width of the computational window that determines the ability to run in parallel, and low number of iterations. The smoothed cutoff method that gradually vanishes the interactions between atoms distant further than cutoff distance should raise the accuracy little up in case that simple cutoff method fails. The method is computationally cheap, so the ratio between evaluation time for $\mathcal{F}$ and $\mathcal{G}$ that determines the speed-up is quite high.

Wolf Summation Method Coarse approximation would use Wolf summation method [WKPE99] to calculate electrostatics. We assume that with little additional computational cost (compared to the cutoff method), we can gain faster convergence of the parareal scheme. The reasons for its suitability remains the same as for *Cutoff Method*, even better convergence is expected.

By analyses of several options for $\mathcal{G}$, we would like to study, implement, and experimentally evaluate especially the approaches based on cheaper methods for calculation of long-range interactions.

3.3 Complexity and ability to run in parallel

Let's assume we are running a simulation with the parareal multilevel summation method with T_W time points in computational window, $T_{total} = WT_W$ time points in whole simulation and K iterations of convergence. For the computation of λ_n^{K-1} for all $n = 1..T_W$, we need $T_W(K+1)-1$ calculations of function $\mathcal{G}$, TK calculations of function $\mathcal{F}$ and TK subtractions to get Δ. Assume we can calculate the function $\mathcal{G}$ with $Q_\mathcal{G}$ floating point

[5] Cutoff distance in MSM is the distance within which the part of potential is calculated on current grid.

operations on $P_{\mathcal{G}}$ processors in time $R_{\mathcal{G}}$ and the function $\mathcal{F}$ with $Q_{\mathcal{F}}$ floating point operations on $P_{\mathcal{F}}$ processors in time $R_{\mathcal{F}}$. If the number of processors depends only on spatial decomposition of calculation, then $P_{\mathcal{G}} \approx P_{\mathcal{F}}$ as we are simulating exactly the same system. For simplicity, let's assume that the overhead of calculation caused by synchronization and awaiting communication should be similar both for $\mathcal{G}$ and $\mathcal{F}$. Therefore, we can assume the ratio of two functions to be

$$\frac{R_{\mathcal{F}}}{R_{\mathcal{G}}} \approx \frac{Q_{\mathcal{F}}}{Q_{\mathcal{G}}} = Q_{\mathcal{F}/\mathcal{G}} \tag{19}$$

The cost of the calculation heavily depends on the distribution of tasks. In this context, a task represents the calculation of the function $\mathcal{G}$ or $\mathcal{F}$ on given input data λ and Δ. The distribution of tasks assigns the tasks to different processors. In preliminary analysis we came up with two distribution plans. Although they differ in number of used processors, the time speed-up compared to the sequential algorithm is directly proportional to $Q_{\mathcal{F}/\mathcal{G}}$ in both.

Calculation distribution plan 1 The simplest plan assigns to computation unit P1[6] calculations of $\mathcal{G}$ for all time points in computational window whereas $T_W = Q_{\mathcal{F}/\mathcal{G}}$. That should take $Q_{\mathcal{F}/\mathcal{G}} R_{\mathcal{G}} = R_{\mathcal{F}}$ time. After that, $Q_{\mathcal{F}/\mathcal{G}}$ units calculate in parallel $\mathcal{F}$ for all points in the computational window. After that, unit P1 calculates $\mathcal{G}$ for subsequent time points in the following computational window and continues analogously. We need only $Q_{\mathcal{F}/\mathcal{G}}$ units. The speed-up is

$$\frac{T_{total} R_{\mathcal{F}}}{2\frac{T_{total}}{Q_{\mathcal{F}/\mathcal{G}}} R_{\mathcal{F}}} = \frac{Q_{\mathcal{F}/\mathcal{G}}}{2} \tag{20}$$

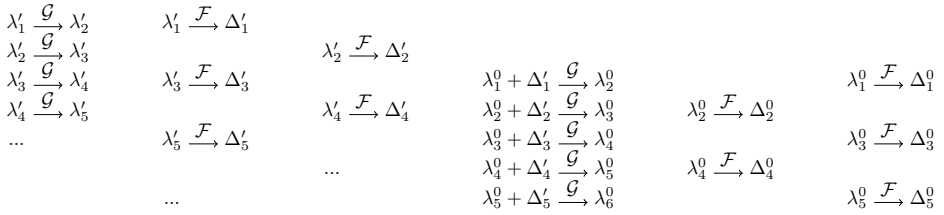
Calculation distribution plan 2 Figure 3.1 depicts a more pipelined plan. It assigns the task of $\mathcal{F}$ calculation to the available (or new) unit as soon as it has the prerequisite—result of $\mathcal{G}$. We need K units for $\mathcal{G}$ tasks and $KQ_{\mathcal{F}/\mathcal{G}}$ units for parallel calculation of $\mathcal{F}$. Time speed-up is

$$\frac{T R_{\mathcal{F}}}{(\frac{T}{Q_{\mathcal{F}/\mathcal{G}}} + K) R_{\mathcal{F}}} = \frac{Q_{\mathcal{F}/\mathcal{G}}}{1 + \frac{K}{T Q_{\mathcal{F}/\mathcal{G}}}} \tag{21}$$

The value of $\frac{K}{T Q_{\mathcal{F}/\mathcal{G}}}$ would be close to 0 in case of the method with high speed-up (proportional to $Q_{\mathcal{F}/\mathcal{G}}$) and good convergence (small K) applied to long simulation time (high T).

[6] In these distribution plans, we consider a computation unit to have $P_{\mathcal{G}}$ or $P_{\mathcal{F}}$ processors.

Figure 3.1: Calculation distribution plan 2.



Example Let's assume $\mathcal{G}$ is simple cutoff method with cutoff 12 Å; $\mathcal{F}$ is MSM with $a = 12$ Å and $h = 2$ Å. Then for $h^* = 1$ Å, $m = 2$ (for C^2 Taylor smoothing function), $p = 3$ (Hermite cubic interpolant)

$$\begin{aligned} Q_{\mathcal{G}} &= 2311N \\ Q_{\mathcal{F}} &= 136813N \\ Q_{\mathcal{F}/\mathcal{G}} &\approx 60 \end{aligned} \tag{22}$$

Two calculation distribution plans presented above have time speed-up proportional to $Q_{\mathcal{F}/\mathcal{G}}$. With appropriately chosen functions we believe we can achieve the speed-up by an order of magnitude. Further distribution plans will be analyzed.

3.4 Convergence

The convergence of the parareal MSM method, i.e. how high the number of iterations is needed to satisfy the acceptance criteria, influences the speed-up and ability to run in parallel in indirectly proportional manner. In case of low number of iterations needed, the width of the computational window can grow and vice versa. The convergence depends mainly on the accuracy and stability of $\mathcal{G}$. Both cutoff method and Wolf summation method are considered rather accurate [BAD05, WKPE99], the simulations do not blow up just because of their errors. Although they do not achieve the accuracy of more sophisticated methods, their cost in number of floating point operations is the lowest.

The error Δ, the difference between $\mathcal{F}$ and $\mathcal{G}$, decreases with increasing cutoff. It represents the change in positions of atoms calculated by $\mathcal{F}$ and positions calculated by $\mathcal{G}$. With increasing number of iterations, the error grows exponentially, however, that does not necessarily mean high divergence. Two different sets of atom positions can in fact represent the same system [YSC06, SC05] if their radial distribution functions coincide. On the other hand, the similarity of atom positions from early time points does not necessarily mean good convergence during the whole simulation time.

Therefore, the actual convergence of the method has to be evaluated by computer experiment with a prototype implementation.

3.5 Future work

The presented proposal of the parareal MSM method as parallel-in-time computation of MD simulations offers several opportunities for future work.

Combined with approach of Yu et al. [YSC06], data from previously done coarse-grained simulation could serve as an assessment of changes in system. The difference between simulations in vacuum and solvated systems would probably cause too many failures of their prediction algorithm, however, the data can serve as high-level estimation.

According to our analyses, we anticipate the speed-up by an order of magnitude. Further acceleration can be achieved through combination of spatial and temporal decomposition, already mentioned in [MT05]. This concept has been successfully implemented for gravitational N-body problem in [SRK+12].

Also, we do not neglect the acceleration that GPUs can provide. The parareal scheme in the second distribution plan nicely flows without major load inbalances if the number of iterations is fixed. The ratio between number of floating point operations and required memory transfers makes the implementation on GPU platform worth trying. Moreover, GPUs implement the evaluation of inverse square root (for Euclidean distance between atoms) in hardware so the number of floating point operations needed is lower. That multiplies the ratio $Q_{\mathcal{F}/\mathcal{G}}$ and makes the speedup (compared with GPU calculation sequential-in-time) by two orders of magnitude possible. The number of supposedly saturated GPUs is directly proportional to the width of the computational window.

The parareal scheme combined with MD algorithm could conceal an interesting possibility. If we are able to roughly approximate Δ without full evaluation of expensive $\mathcal{F}$, we can save the calculation when the correction is not necessary. Then, an adaptive control of the simulation could quickly calculate results in phases when the system just moves around the minimum a little bit. The interesting part with major changes would be evaluated by the accurate function step-by-step.

3.6 Evaluation

We will evaluate several aspects of the proposed methods. As many of them require experimental evaluation, we will develop a prototype implementation. With data acquired by computer experiments we will compare accuracy and efficient parallelization with other parallel programs.

First, we will examine the method by feasible simulations of medium size and length at computational resources provided by CERIT-SC and Meta-

centrum NGI, Czech Republic. We will compare the accuracy of the results acquired by the parareal MSM and by sequential-in-time MSM. The accuracy of MSM has been extensively analyzed and compared with other methods in [Har06]. Our comparison should uncover possible issues introduced by time parallelism. Strong scalability analysis can show its potential even with a few hundreds of cores, if the ratio between the number of atoms and number of cores gets close to 0 (as achieved by Copernicus). Also, we will examine the speed-up from classical MSM to parareal MSM to unfold the contribution of time parallelism. Another measure of parallel MD code is the speed in ns/day, i.e. how many ns of simulation time we can compute in 24 hours. However, that value depends on the simulated system and the parameters of the infrastructure, so it needs to be interpreted in that context.

If our approach shows promising results on medium-sized simulations, we will proceed with more elaborate and optimized implementation. The large and long simulation of the virus passing the cell's membrane, infeasible with current methods, would require high number of computational resources. PRACE Research Infrastructure and ScalaLife project offer access to supercomputers in Europe, e.g. to half a million cores in Jülich SuperComputing Centre or the supercomputer with expected performance 1 PFLOP/s that is being built in Ostrava, Czech Republic, by IT for Innovations, National Supercomputing Centre. With such large and long simulation run, we can compare the weak and strong scalability of our approach with the methods mentioned in subsection 2.4. Moreover, as Pavel Plevka, CEITEC, will study this process by experimental means, we plan to compare the results of our simulation to experimentally acquired ones.

4 Aim of the Work

Proposed thesis aims to research how we can simulate large molecular dynamics systems for long simulation times. The acceleration of parallel implementations hits the strong scalability limit. We want to increase the level of parallelism by introducing the concept of parallel-in-time computation to molecular dynamics.

Published solutions with time parallelism have both promising [PLP+11] and disappointing results [BBM+02]. The combination of the parareal method and MD code has not achieved high speed-up or reasonable convergence, however, we suspect that inappropriately chosen coarse function cause that. We will further study and develop more suitable coarse functions. After this thesis, it will be known whether the fine-grained time parallelism without *a priori* knowledge increases the strong scalability of MD algorithm.

We will investigate several hypotheses through extensive study and evaluation of experiments with prototype implementation. We will analyze how different coarse functions affect the convergence and speed-up of the calculation and if the method gives results comparable with sequential-in-time methods in accuracy. Our concern will focus on the strong scalability—how many cores we can saturate with decreasing time to result and how much we can reduce the ratio between the size of the system and the number of resources.

4.1 Expected Results

We expect to contribute to both computational chemistry and computer science. If successful, the time parallelism could, to some extent, solve the sampling problem, and serve as inspiration for other computational chemistry tools to incorporate it into their calculations. Research in high performance computing is more and more interested in algorithms that are able to utilize large computational resources. From the point of decomposition algorithms for parallel and distributed computing, an application of the combined spatial and temporal decomposition shows perspective for other simulations.

We will implement a prototype algorithm to experimentally evaluate the accuracy, convergence, speed-up, and strong scalability on systems with feasible size and simulation time. If the prototype algorithm shows high speed-up and scaling we want to apply for computational time of a supercomputer. A simulation of the virus passing the cell's membrane with possibly hundreds of millions atoms for at least 50 μs would verify properties of the method at large scale.

We plan to publish the method and results in related conferences and scientific journals, such as:

Conferences

- IEEE International Conference on High Performance Computing and Simulation
- ACM/IEEE/SCS Workshop on Parallel and Distributed Simulation
- The International ACM Symposium on High-Performance Parallel and Distributed Computing
- International Conference on High Performance Computing & Simulation
- International Conference on High Performance Computing
- Supercomputing Conference
- Summer Computer Simulation Conference

Scientific Journals

- Journal of Computational Chemistry
- Journal of Computational Physics
- Journal of Chemoinformatics
- Journal of Chemical Information and Modeling
- Journal of Chemical Theory and Computation
- Journal of Chemical Physics
- Computer Physics Communications
- Parallel Algorithms and Applications
- SIAM Journal on Scientific Computing
- Parallel Computing
- Domain Decomposition Methods in Science and Engineering
- Journal of Parallel and Distributed Computing

The text of PhD. thesis will, as common, analyze the field, state an interesting and difficult problem, propose a novel solution and evaluate it.

4.2 Schedule

We schedule future progress with the following plan.

- 01/2014 - 05/2014: implementation of the prototype MD code with MSM for evaluation of long-range interactions
- 06/2014 - 07/2014: evaluation of the convergence of the parareal scheme, conference article writing
- 06/2014 - 11/2014: implementation of the parareal scheme with MSM as the fine function and cutoff/Wolf summation method as coarse functions
- 12/2014 - 01/2015: evaluation of the speed-up and strong scalability of the parareal scheme on feasible simulations, journal article writing
- 01/2015 - 06/2015: optimizations of the prototype
- 06/2015 - 12/2015: implementation of GPU acceleration, conference article writing, applying for computing time for the simulation of the virus
- 01/2016 - 06/2016: thesis writing, the simulation run
- 06/2016: submission of the thesis
- 06/2016 - 12/2016: journal article writing, postprocessing data from the simulation run
- 12/2016: the thesis defense

Glossary

computation time the sum of time durations of the algorithm run as measured by clocks of all processors that took part.

computational chemistry a field of chemistry that uses computational resources to model approximative representations of chemical systems.

computer experiment a series of runs of the algorithm that provide data for evaluation of algorithm properties.

conformation a distinct 3D structure of a protein in a local minimum of the energy surface.

energy surface a function that describes the energy of the system with respect to coordinates of all atoms.

exascale computing computing systems achieving the performace of exaFLOPs, according to predictions exascale computing system will be built till 2020.

experiment an experiment in sense of traditional experimental chemistry.

explicit model of solvent the solvent is represented by explicit atoms of solvent molecules that take part in calculation in the same way as atoms of the molecule of interest.

force field a set of parameters and equations that empirically describe the interactions of atoms.

fs femtosecond, 10^{-15} s.

implicit model of solvent the solvent is not explicitly represented by its molecules, its effects are included in the parameters of the system.

in sillico experiment an experiment done by computer simulation.

molecular dynamics a method of computational chemistry that calculates the movements of atoms caused by their interactions; this work deals with molecular dynamics with molecular mechanics model.

molecular mechanics a model of computational chemistry that considers atoms as undivided points of mass, it describes their interactions with empirical functions, it does not explicitly account for electrons or quantuum effects.

molecule of interest a molecule of biological or chemical interest, e.g. a protein, a virus, a nanotube.

protein folding a biochemical process in which a protein folds, i.e. changes its 3D structure in fluent movement; commonly simulated by molecular dynamics.

sampling problem the energy surface usually has many local minima split by energy barriers that are crossed only with some probability; it can take quite long for molecular dynamics simulation to cover all minima, i.e. to sample the space of energy surface, hence the sampling problem.

simulation time the time duration of simulated process; often number of steps $\times$ length of single timestep.

solvent usually molecules of water and salt ions that surround the molecule of interest in the simulation.

strong scaling a function that maps increasing number of computational resources that solve the problem of fixed size to the wallclock time of the computation; it shows how the wallclock time (usually) decreases as the resources grow but the problem remains the same.

wallclock time the time duration of the algorithm run as measured by wall clock.

weak scaling a function that maps increasing number of computational resources that solve the problem of increasing size to the wallclock time of the computation; it shows how the wallclock time ideally remains the same as the resources grow and the problem grows.

Acronyms

Å angstrom, 10^{-10} m, 0.1 nm.

exa 10^{18}.

FLOP floating point operation.

FLOPs floating point operations per second, a unit for measurement of performance of computing systems.

fs femtosecond, 10^{-15} second.

MD molecular dynamics, in this work with molecular mechanics model.

MM molecular mechanics.

MM MD molecular dynamics with molecular mechanics model of interactions.

MSM multilevel summation method, a method for calculation of long-range interactions [Har06].

Bibliography

[ABF+13] A Arnold, M Bolten, F Fahrenberger, F Dachsel, F Gähler, R Halver, F Heber, M Hofmann, J Iseringhausen, I Kabadshow, O Lenz, M Pippig, D Potts, and G Sutmann. A Comparison of Scalable Fast Methods for Long-Range Interactions. *Physical Review E*, 88(6):063308:1–063308:22, 2013.

[ALT08] JA Anderson, CD Lorenz, and A Travesset. General purpose molecular dynamics simulations fully implemented on graphics processing units. *Journal of Computational Physics*, 227(10):5342–5359, 2008.

[And83] H Andersen. Rattle: A "velocity" version of the shake algorithm for molecular dynamics calculations. *Journal of Computational Physics*, 52(1):24–34, 1983.

[AT89] MP Allen and DJ Tildesley. *Computer simulation of liquids*. Oxford University Press, New York, 1989.

[Aub11] E Aubanel. Scheduling of tasks in the parareal algorithm. *Parallel Computing*, 37(3):172–182, 2011.

[AW59] BJ Alder and TE Wainwright. Studies in Molecular Dynamics. I. General Method. *Journal of Chemical Physics*, 31(2):459–466, 1959.

[AYF+13] Y Andoh, N Yoshii, K Fujimoto, K Mizutani, H Kojima, A Yamada, S Okazaki, K Kawaguchi, H Nagao, K Iwahashi, F Mizutani, K Minami, S Ichikawa, H Komatsu, S Ishizuki, Y Takeda, and M Fukushima. MODYLAS: A Highly Parallelized General-Purpose Molecular Dynamics Simulation Program for Large-Scale Systems with Long-Range Forces Calculated by Fast Multipole Method (FMM) and Highly Scalable Fine-Grained New Parallel Processing Algorithms. *Journal of Chemical Theory and Computation*, 9(7):3201–3209, 2013.

[BAD05] DAC Beck, RS Armen, and V Dagget. Cutoff Size Need Not Strongly Influence Molecular Dynamics results for Solvated Polypeptides. *Biochemistry*, 44(2):609–616, 2005.

[Bak05] NA Baker. Improving implicit solvent simulations: a Poisson-centric view. *Current Opinion in Structural Biology*, 15(2):137–143, 2005.

[Bal05] G Bal. On the Convergence and the Stability of the Parareal Algorithm to solve Partial Differential Equations. In *Domain Decomposition Methods in Science and Engineering*, volume 0, chapter XI, pages 425–432. Springer Berlin Heidelberg, 2005.

[BBM+02] L Baffico, S Bernard, Y Maday, G Turinici, and G Zérah. Parallel-in-time molecular-dynamics simulations. *Physical Review E*, 66:057701:1–057701:4, 2002.

[BBSD13] L Baweja, K Balamurugan, V Subramanian, and A Dhawan. Hydration patterns of graphene-based nanomaterials (GBNMs) play a major role in the stability of a helical protein: a molecular dynamics simulation study. *ACS Journal of Surfaces and Colloids*, 29(46):14230–14238, 2013.

[BdOD+13] L Boechi, CAF de Oliveira, I Da Fonseca, K Kizjakina, P Sobrado, JJ Tanner, and JA McCammon. Substrate-dependent dynamics of UDP-galactopyranose mutase: Implications for drug design. *Protein Science*, 22(11):1490–1501, 2013.

[BDS05] KJ Bowers, RO Dror, and DE Shaw. Overview of neutral territory methods for the parallel evaluation of pairwise particle interactions. *Journal of Physics: Conference Series*, 16:300–304, 2005.

[BFH+] M Bolten, F Fahrenberger, R Halver, F Heber, M Hofmann, I Kabadshow, O Lenz, M Pippig, and G Sutmann. ScaFaCoS, C subroutine library. http://scafacos.github.com/.

[BH86] J Barnes and P Hut. A hierarchical O($N \log N$) force-calculation algorithm. *Nature*, 324:446–449, 1986.

[BHN03] H Berman, K Henrick, and H Nakamura. Announcing the worldwide Protein Data Bank. *Nature Structural Biology*, 10(12):980–981, 2003. http://www.wwpdb.org.

[BLMM13] A Baudron, J Lautard, Y Maday, and O Mula. The parareal in time algorithm applied to the kinetic neutron diffusion

equation. In *International Conference on Domain Decomposition Methods*, 2013.

[BM02] G Bal and Y Maday. A "parareal" time discretization for non-linear pde's with application to the pricing of an american put. In *Recent Developments in Domain Decomposition Methods*, volume 23, pages 189–202. 2002.

[Bul13] J Bulin. Large-scale time parallelization for molecular dynamics problems. Technical report, Royal Institute Of Technology, Stockholm, Stockholm, 2013.

[BvHM00] WL Briggs, E van Henson, and SF McCormick. *A Multigrid Tutorial*. SIAM Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2nd edition, 2000.

[CCB+95] WD Cornell, P Cieplak, CI Bayly, IR Gould, KM Merz, DM Ferguson, DC Spellmeyer, T Fox, JW Caldwell, and PA Kollman. A Second Generation Force Field for the Simulation of Proteins, Nucleic Acids, and Organic Molecules. *Journal of the American Chemical Society*, 117(19):5179–5197, 1995.

[CCD+05] DA Case, TE Cheatham, T Darden, H Gohlke, R Luo, KM Merz, A Onufriev, C Simmerling, B Wang, and RJ Woods. The Amber biomolecular simulation programs. *Journal of computational chemistry*, 26(16):1668–1688, 2005.

[CHB+05] M Christen, PH Hünenberger, D Bakowies, R Baron, R Bürgi, DP Geerke, TN Heinz, MA Kastenholz, V Kräutler, C Oostenbrink, C Peter, D Trzesniak, and WF van Gunsteren. The GROMOS software for biomolecular simulation: GROMOS05. *Journal of computational chemistry*, 26(16):1719–1751, 2005.

[DH98] M Deserno and C Holm. How to mesh up Ewald sums. I. A theoretical and numerical comparison of various particle mesh routines. *Journal of Chemical Physics*, 109(18):7678–7693, 1998.

[DMDS95] A Deshpande, S Malhotra, CC Douglas, and MH Schultz. A rigorous analysis of time domain parallelism. *Parallel Algorithms and Applications*, 6(1):53–62, 1995.

[DYP93] T Darden, DM York, and L Pedersen. Particle Mesh Ewald: An $N \log(N)$ method for Ewald sums in large systems. *Journal of Chemical Physics*, 99:10089–10092, 1993.

[EPB+95] U Essmann, L Perera, ML Berkowitz, T Darden, H Lee, LG Pedersen, and Et al. A smooth particle mesh Ewald method. *Journal of Chemical Physics*, 103(November):31–34, 1995.

[Ewa21] PP Ewald. Die Berechnung optischer und elektrostatischer Gitterpotentiale. *Annalen der Physik*, 369(3):253–287, 1921. cited by [DH98].

[GaOdC+11] AP Guimarães, AA Oliveira, EFF da Cunha, TC Ramalho, and TCC França. Design of new chemotherapeutics against the deadly anthrax disease. Docking and molecular dynamics studies of inhibitors containing pyrrolidine and riboamidrazone rings on nucleoside hydrolase from Bacillus anthracis. *Journal of Biomolecular Structure & Dynamics*, 28(4):455–469, 2011.

[Gea71] CW Gear. The automatic integration of ordinary differential equations. *Communications of the ACM*, 14(3):176–179, 1971.

[GR87] L Greengard and V Rokhlin. A Fast Algorithm for Particle Simulations. *Journal of Computational Physics*, 73:325–348, 1987.

[Gro13] Gromacs. Blowing up. http://www.gromacs.org/Documentation/Terminology/Blowing_Up, 2013.

[GT98] H Grubmueller and P Tavan. Multiple Time Step Algorithms for Molecular Dynamics Simulations of Proteins: How Good Are They? *Journal of Computational Chemistry*, 19(13):1534–1552, 1998.

[GV07] MJ Gander and S Vandewalle. Analysis of the Parareal Time-Parallel Time-Integration Method. *SIAM Journal on Scientific Computing*, 29(2):556–578, 2007.

[Hal96] TA Halgren. Merck Molecular Force Field. I. Basis, Form, Scope, Parametrization, and Performance of MMFF94. *Journal of Computational Chemistry*, 17:490–519, 1996.

[Har06] DJ Hardy. *Multilevel summation for the fast evaluation of forces for the simulation of biomolecules*. PhD thesis, University of Illinois at Urbana-Champaign, 2006.

[HBBF97] B Hess, H Bekker, HJC Berendsen, and JGEM Fraaije. LINCS: A linear constraint solver for molecular simulations. *Journal of Computational Chemistry*, 18(12):1463–1472, 1997.

[HE88] RW Hockney and JW Eastwood. *Computer simulation using particles*. Taylor & Francis, Inc., Bristol, PA, USA, 1988.

[Hes08] B Hess. P-LINCS: A Parallel Linear Constraint Solver for Molecular Simulation. *Journal of Chemical Theory and Computation*, 4(1):116–122, 2008.

[HKvdSL08] B Hess, C Kutzner, D van der Spoel, and E Lindahl. GROMACS 4: Algorithms for Highly Efficient, Load-Balanced, and Scalable Molecular Simulation. *Journal of Chemical Theory and Computation*, 4(3):435–447, 2008.

[Hor92] G Horton. The time-parallel Multigrid Method. *Communications in Applied Numerical Methods*, 8:585–595, 1992.

[HSP+04] HW Horn, WC Swope, JW Pitera, JD Madura, TJ Dick, GL Hura, and T Head-Gordon. Development of an improved four-site water model for biomolecular simulations: TIP4P-Ew. *Journal of Chemical Physics*, 120(20):9665–9678, 2004.

[HSS09] DJ Hardy, JE Stone, and K Schulten. Multilevel Summation of Electrostatic Potentials Using Graphics Processing Units. *Parallel computing*, 35(3):164–177, 2009.

[HV95] G Horton and S Vandewalle. A space-time multigrid method for parabolic partial differential equations. *SIAM Journal on Scientific Computing*, 16(4):848–864, 1995.

[Jen07] F Jensen. *Introduction to computational chemistry*. John Wiley & Sons Ltd, Great Britain, 2 edition, 2007.

[KM02] M Karplus and JA McCammon. Molecular dynamics simulations of biomolecules. *Nature Structural Biology*, 9(9):646–652, 2002.

[Koe06] P Koehl. Electrostatics calculations: latest methodological advances. *Current opinion in structural biology*, 16(2):142–151, 2006.

[LB89] RJ Loncharich and BR Brooks. The effects of truncating long-range forces on protein dynamics. *Proteins*, 6(1):32–45, 1989.

[Lea01] AR Leach. *Molecular modelling: principles and applications*. Pearson Education, Dorchester, 2nd edition, 2001.

[Lew10] E Lewars. *Computational Chemistry: Introduction to the Theory and Applications of Molecular and Quantum Mechanics*. Springer, 2nd edition, 2010.

[LG08] A Laio and FL Gervasio. Metadynamics: a method to simulate rare events and reconstruct the free energy in biophysics, chemistry and material science. *Reports on Progress in Physics*, 71(12):126601:1–126601:22, 2008.

[LH11] C Lee and S Ham. Characterizing amyloid-beta protein misfolding from molecular dynamics simulations with explicit water. *Journal of Computational Chemistry*, 32(2):349–355, 2011.

[LL12] D Lau and R Lam. Atomistic Prediction of Nanomaterials: Introduction to Molecular Dynamics Simulation and a Case Study of Graphene Wettability. *IEEE Nanotechnology Magazine*, 6(1):8–13, 2012.

[LMT01] JL Lions, Y Maday, and G Turinici. Résolution d'EDP par un schéma en temps «pararéel ». *Comptes Rendus de l'Académie des Sciences - Series I - Mathematics*, 332(7):661–668, 2001. cited by [Mad08].

[LRSV82] E Lelarasmee, AE Ruehli, and AL Sangiovanni-Vincentelli. The Waveform Relaxation Method for Time-Domain Analysis of Large Scale Integrated Circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 1(3):131–145, 1982.

[LSOS04] MS Lee, FR Salsbury Jr., MA Olson, and FR Salsbury. An efficient hybrid explicit/implicit solvent method for biomolecular simulations. *Journal of Computational Chemistry*, 25(16):1967–1978, 2004.

[LSSP02] SM Larson, CD Snow, M Shirts, and VS Pande. Folding@Home and Genome@Home: Using distributed computing to tackle previously intractable problems in computational biology. Technical report, Stanford University, 2002. http://arxiv.org/pdf/0901.0866v1.pdf.

[LSVMW08] W Liu, B Schmidt, G Voss, and W Müller-Wittig. Accelerating molecular dynamics simulations using Graphics Processing Units with CUDA. *Computer Physics Communications*, 179(9):634–641, 2008.

[Mad08] Y Maday. The parareal in time algorithm. *Technical Report R08030, Université Pierre et Marie Curie*, pages 1–24, 2008. https://www.ljll.math.upmc.fr/publications/2008/R08030.pdf.

[Mat13] C Matthews. *The error in the invariant measure of numerical discretization schemes for canonical sampling of molecular dynamics*. PhD thesis, University of Edinburgh, 2013.

[MBF01] DA MacKerell, N Banavali, and N Foloppe. Development and current status of the CHARMM force field for nucleic acids. *Biopolymers*, 56(4):257–265, 2001.

[MF04] R Miron and K Fichthorn. Multiple-Time Scale Accelerated Molecular Dynamics: Addressing the Small-Barrier Problem. *Physical Review Letters*, 93(12):128301:1–128301:4, 2004.

[MK92] S Miyamoto and PA Kollman. Settle: An analytical version of the SHAKE and RATTLE algorithm for rigid water models. *Journal of Computational Chemistry*, 13(8):952–962, 1992.

[ML67] WL Miranker and W Liniger. Parallel methods for the numerical integration of ordinary differential equations. *Mathematics of Computation*, 21(99):303–320, 1967.

[MSH07] B Markvoort, P Spijker, and P Hilbers. *Theory explanation to PumMa tool*. BioModeling and bioInformatics group of the Department of Biomedical Engineering at the Eindhoven University of Technology, 2007.

[MT05] Y Maday and G Turinici. The Parareal in Time Iterative Solver: a Further Direction to Parallel Implementation. In *Domain Decomposition Methods in Science and Engineering*, pages 441–448, 2005.

[Nie64] J Nievergelt. Parallel methods for integrating ordinary differential equations. *Communications of the ACM*, 7(12):731–733, 1964.

[NVK93] A Nakano, P Vashishta, and RK Kalia. Parallel multiple-time-step molecular dynamics with three-body interaction. *Computer Physics Communications*, 77:303–312, 1993.

[OCP+13] J Ostmeyer, S Chakrapani, AC Pan, E Perozo, and B Roux. Recovery from slow inactivation in K+ channels is controlled by water molecules. *Nature*, 501(7465):121–124, 2013.

[PBW+05] JC Phillips, R Braun, W Wang, J Gumbart, E Tajkhorshid, E Villa, C Chipot, RD Skeel, L Kalé, and K Schulten. Scalable molecular dynamics with NAMD. *Journal of Computational Chemistry*, 26(16):1781–1802, 2005.

[PCC+95] DA Pearlman, DA Case, JW Caldwell, WS Ross, TE Cheatham, S DeBolt, DM Ferguson, G Seibel, and P Kollman. AMBER, a package of computer programs for applying molecular mechanics, normal mode analysis, molecular dynamics and free energy calculations to simulate the structural and energetic properties of molecules. *Computer Physics Communications*, 91(1-3):1–41, 1995.

[PDD11] EA Proctor, F Ding, and NV Dokholyan. Discrete molecular dynamics. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 1(1):80–92, 2011.

[Pli95] SJ Plimpton. Fast Parallel Algorithms for Short- Range Molecular Dynamics. *Journal of Computational Physics*, 117(June 1994):1–42, 1995.

[Pli03] SJ Plimpton. LAMMPS web page, 2003. http://lammps.sandia.gov.

[PLP+11] S Pronk, P Larsson, I Pouya, GR Bowman, IS Haque, K Beauchamp, B Hess, VS Pande, PM Kasson, and E Lindahl. Copernicus: a new paradigm for parallel adaptive molecular dynamics. In *Proceedings of Supercomputing*, SC '11, pages 60:1–60:10, New York, NY, USA, 2011. ACM.

[RB98] R Rudd and J Broughton. Coarse-grained molecular dynamics and the atomic limit of finite elements. *Physical Review B*, 58(10):R5893–R5896, 1998.

[RCB77] JP Ryckaert, G Ciccotti, and HJC Berendsen. Numerical integration of the cartesian equations of motion of a system with constraints: molecular dynamics of n-alkanes. *Journal of Computational Physics*, 23(3):327–341, 1977.

[RCC+92] AK Rappé, CJ Casewit, KS Colwell, WA Goddard, and WM Skiff. UFF, a Full Periodic Table Force Field for Molecular Mechanics and Molecular Dynamics Simulations. *Journal of American Chemical Society*, 2(114):10024–10035, 1992.

[RGC+09] DF Richards, JN Glosli, B Chan, MR Dorr, EW Draeger, JL Fattebert, WD Krauss, T Spelce, FH Streitz, MP Surh, and JA Gunnels. Beyond Homogeneous Decomposition Scaling

Long-Range Forces on Massively Parallel Systems. In *Proceedings Of Supercomputing*, SC '09, pages 60:1–60:12, New York, NY, USA, 2009. ACM.

[Ric04] SW Rick. A reoptimization of the five-site water potential (TIP5P) for use with Ewald sums. *Journal of Chemical Physics*, 120(13):6085–6093, 2004.

[SABK82] WC Swope, HC Andersen, PH Berens, and Wilson KR. *Journal of Chemical Physics*, (76):637–648, 1982.

[SC05] A Srinivasan and N Chandra. Latency tolerance through parallelization of time in scientific applications. *Parallel Computing*, 31(7):777–796, 2005.

[SD01] C Sagui and T Darden. Multigrid methods for classical molecular dynamics simulations of biomolecules. *Journal of Chemical Physics*, 114:6578–6591, 2001.

[SDD+07] DE Shaw, MM Deneroff, RO Dror, JS Kuskin, RH Larson, JK Salmon, C Young, B Batson, KJ Bowers, JC Chao, MP Eastwood, J Gagliardo, JP Grossman, RC Ho, DJ Ierardi, I Kolossváry, JL Klepeis, T Layman, C McLeavey, MA Moraes, R Mueller, EC Priest, Y Shan, J Spengler, M Theobald, B Towles, and SC Wang. Anton, a special-purpose machine for molecular dynamics simulation. *Proceedings of Annual International Symposium on Computer Architecture*, 35(2):1–12, 2007.

[SKE+05] Y Shan, JL Klepeis, MP Eastwood, RO Dror, and DE Shaw. Gaussian split Ewald: A fast Ewald mesh method for molecular simulation. *Journal of Chemical Physics*, 122:054101:1–054101:13, 2005.

[SKS11] B Schaefer, C Kisker, and CA Sotriffer. Molecular dynamics of Mycobacterium tuberculosis KasA: implications for inhibitor and substrate binding and consequences for drug design. *Journal of Computer-aided Molecular Design*, 25(11):1053–1069, 2011.

[SLPS09] R Schultz, B Lindner, L Petridis, and JC Smith. Scaling of Multimillion-Atom Biological Molecular Dynamics Simulation on a Petascale Supercomputer. *Journal of Chemical Theory and Computation*, 5(10):2798–2808, 2009.

[SNPG02] CD Snow, H Nguyen, Vijay S Pande, and Martin Gruebele. Absolute comparison of simulated and experimental protein-folding dynamics. *Nature*, 420(6911):102–106, 2002.

[SO99] Y Sugita and Y Okamoto. Replica exchange molecular dynamics method for protein folding. *Chemical Physics Letters*, 314(November):141–151, 1999.

[SPF+07] JE Stone, JC Phillips, PL Freddolino, DJ Hardy, LG Trabuco, and K Schulten. Accelerating molecular modeling applications with graphics processors. *Journal of Computational Chemistry*, 28(16):2618–2640, 2007.

[SRK+12] R Speck, D Ruprecht, R Krause, M Emmett, M Minion, M Winkel, and P Gibbon. A massively space-time parallel N-body solver. In *Proceedings of Supercomputing*, pages 92:1–92:11, 2012.

[SRn03] GA Staff and EM Rø nquist. Stability of the parareal algorithm. In *Proceedings of International Domain Decomposition Conference*, pages 449–456. Springer, 2003.

[SS92] H Schreiber and O Steinhauser. Cutoff size does strongly influence molecular dynamics results on solvated polypeptides. *Biochemistry*, 31(25):5856–5860, 1992.

[SS05] G Sutmann and B Steffen. A particle-particle particle-multigrid method for long-range interactions in molecular simulations. *Computer Physics Communications*, 169(1-3):343–346, 2005.

[SS07] M Sotomayor and K Schulten. Single-molecule experiments in vitro and in silico. *Science*, 316(5828):1144–1148, 2007.

[ST82] K Stuben and U Trottenberg. *Multigrid methods: Fundamental algorithms, model problem analysis and applications*, volume 960. Springer Berlin Heidelberg, Berlin, 1982.

[ST05] RD Skeel and PF Tupper. Mathematical Issues in Molecular Dynamics. *Banff International Research Station Reports*, 2005.

[STH02] RD Skeel, I Tezcan, and DJ Hardy. Multiple grid methods for classical molecular dynamics. *Journal of computational chemistry*, 23(6):673–84, 2002.

[TP10] SS Tallury and MA Pasquinelli. Molecular dynamics simulations of flexible polymer chains wrapping single-walled carbon nanotubes. *Journal of Physical Chemistry B*, 114(12):4122–4129, 2010.

[TR85] K Toukan and A Rahman. Molecular-dynamics study of atomic motions in water. *Physical Review B*, 31(5):2643–2648, 1985.

[TV77] GM Torrie and JP Valleau. Nonphysical Sampling Distributions in Monte Carlo Free-Energy Estimation: Umbrella Sampling. *Journal of Computational Physics*, 23:187–199, 1977.

[vDDLG01] ACT van Duin, S Dasgupta, F Lorant, and WA Goddard. ReaxFF: A Reactive Force Field for Hydrocarbons. *Journal of Physical Chemistry A*, 105(41):9396–9409, 2001.

[VM05] CL Vizcarra and SL Mayo. Electrostatics in computational protein design. *Current Opinion in Structural Biology*, 9(6):622–626, 2005.

[VMVF11] R Vácha, FJ Martinez-Veracoechea, and D Frenkel. Receptor-Mediated Endocytosis of Nanoparticles of Various Shapes. *Nano Letters*, 11(12):5391–5395, 2011.

[VvdV94] S Vandewalle and E van de Velde. Space-time concurrent multigrid waveform relaxation. *Annals of Numerical Mathematics*, 1(1-4):335–346, 1994.

[WF06] H Waisman and J Fish. A space – time multilevel method for molecular dynamics simulations. *Computational Methods in Applied Mechanics and Engineering*, 195:6542–6559, 2006.

[WKPE99] D Wolf, P Keblinski, SR Phillpot, and J Eggebrecht. Exact method for the simulation of Coulombic systems by spherically truncated, pairwise r^{-1} summation. *Journal of Chemical Physics*, 110(17):8254–8282, 1999.

[YSC06] Y Yu, A Srinivasan, and N Chandra. Scalable Time-Parallelization of Molecular Dynamics Simulations in Nano Mechanics. In *Conference on Parallel Processing*, pages 119–126. Ieee, 2006.

[ZPY+13] G Zhao, JR Perilla, EL Yufenyuy, X Meng, B Chen, J Ning, J Ahn, AM Gronenborn, K Schulten, C Aiken, and P Zhang. Mature HIV-1 capsid structure by cryo-electron microscopy and all-atom molecular dynamics. *Nature*, 497(7451):643–646, 2013.